

Approximating Functions on a Mesh with Restricted Voronoi Diagrams

V. Nivolières¹ and B. Lévy²

¹ACE group, University of Liège, Belgium

²ALICE Group, INRIA Grand Est, LORIA, Nancy, France

Abstract

We propose a method that computes a piecewise constant approximation of a function defined on a mesh. The approximation is associated with the cells of a restricted Voronoi diagram. Our method optimizes an objective function measuring the quality of the approximation. This objective function depends on the placement of the samples that define the restricted Voronoi diagram and their associated function values. We study the continuity of the objective function, derive the closed-form expression of its derivatives and use them to design a numerical solution mechanism. The method can be applied to a function that has discontinuities, and the result aligns the boundaries of the Voronoi cells with the discontinuities. Some examples are shown, suggesting potential applications in image vectorization and compact representation of lighting.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

1. Introduction

In this paper, we introduce a new method to compute an approximation of a function on a 3d mesh. Sampling functions is a problem that appears in many different fields and has lots of applications, such as storing attributes on meshes. Our approximation uses a set of samples located in 3d space around the mesh. With the restricted Voronoi diagram of the samples, the mesh is segmented into regions in which the approximation takes a constant value. An example is shown on [Figure 1](#). Our goal is to optimize both the positions of the samples and the associated approximation values. This approach offers several advantages :

- the approximation is not stored directly on the combinatorics of the mesh. The surface can be remeshed without impacting the approximation built on it ;
- the size of the storage is fully controlled by the number of samples, and therefore by the user.
- the approximated function is used as a *black box*. We here mean that we solely need to evaluate the *value* of the function at any point on the surface (no gradient required) ;
- it fully automatically captures the sharp features of the approximated function, without any heuristic to locate these features explicitly ;



Figure 1: Direct light sampling. Left, the original lighting. Center, our approximation with 1000 samples. Right, optimized restricted Voronoi diagram.

We formulate this problem as the minimization of an objective function ([Section 4.2](#)). Using Reynolds transport theorem, we provide a closed form expression for its derivative ([Section 4.5](#)), and study its continuity ([Section 4.6](#)). An iterative procedure is then provided to minimize our objective function ([Section 5](#)). In [Section 6](#), we apply our method to image vectorization and compact representation of lighting on meshes.

2. Background and related work

2.1. Segmentation

Segmentation is one of the most obvious applications for our work. In this context, the goal is to split an image or a mesh into a set of regions, depending on a attribute such as color. Hoiem et al. [HEH05] use the technique designed by Felzenszwalb [FH04] to build pixel clusters in an image. This method performs the segmentation using a greedy graph clustering. Compared to our approach, the resulting regions are not restricted to convex shapes. This technique strongly uses the fact that the image is described by a set of pixels with a grid structures, and therefore cannot handle f as a black box as we do. A similar clustering technique is used by Lecot and Lévy [LL06] in the context of image vectorization.

On 3d meshes, the method proposed by Cohen-Steiner et al. [CSAD04] shares a lot of similarities with ours. The main difference is that their technique applies on dense meshes, and clusters triangles of the mesh, whereas ours uses a restricted Voronoi diagram. We therefore can handle very coarse meshes, at the cost of topology issues when the surface is thin. As shown on Figure 9, these issues do not weaken our results. Another difference is that the clustering and the choice of a representative for each cluster are done in separate passes in this work, whereas we integrate our choice of a representative (the approximation values) in the objective function (Section 4.2) and fully describe the problem with a single objective function.

2.2. Image vectorization

As shown in Section 6.2, our method can be applied as an image vectorization technique. We must however admit that our current choice for a piecewise constant approximation makes our results visually less pleasing than the latest research in the field, which is more focused on describing smooth gradients [LL06, OBW*08, SLWS07]. The main advantage we have over these methods is that we do not need the approximated function f to be described with a set of pixels, and only require to be able to evaluate f at any point of the domain. This allows us to approximate functions on more general domains like meshes.

2.3. Light sampling on surfaces

Lehtinen et al. [LZT*08] describe a hierarchical sampling method for storing light on meshes. Their representation does not depend on the combinatorics of the underlying mesh. This approach is in essence different from ours, since their goal is to effectively compute the lighting on the sampling, whereas we aim at generating the sampling from a known function. Because of its hierarchical structure, this approach is able to consider light interaction at different scales between the objects of the scene. In contrast, our

method adapts the density of the sampling to the variations of the approximated function, and captures particularly well the discontinuities such as shadow boundaries.

Discontinuities require a special attention. Holzschuch and Alonso [HA*04] use the notion of *discontinuity meshing* to capture the shadows in a radiosity simulation. Various approaches have been developed to segment a mesh along the discontinuities in the lighting. Heckbert [Hec92] and Litschinski et al. [LTG92] provide an algorithm to cut a mesh along shadow boundaries. Durand et al. [PV93] offer a definition and an algorithm to compute the *visibility complex*, which describes the visibility between any pair of points in the scene. The complexity of these algorithms rises rapidly with the complexity of the scene. Our method does not provide exact shadow boundary information, but computes the best approximation of the discontinuities given a fixed number of samples. This is done without any particular discontinuity detection mechanism.

2.4. Centroidal Voronoi tessellations

The objective function we define and minimize in this work is a generalization of the objective functions minimized for the generation of centroidal Voronoi tessellations. These techniques usually aim at the generation of regular samplings. Okabe and Suzuki [OS97] list many applications related to the minimization of such functions. The corresponding objective function is of class C^2 [LWL*09]. More recently, Lévy and Bonneel [BB*12] use such a minimization for the anisotropic remeshing of surfaces, while de Goes et al. [dGBOD12] optimize *power diagrams* for stippling applications. Because it aims at sampling any function provided by the user, our method is very similar to this last approach. The main difference is that our approximation is defined by approximation values associated with each sample, whereas the work of de Goes et al. approximate the function with the density of the samples. Their method therefore increases the number of samples where the value of the approximated function is high, whereas our number of samples is provided by the user, and these samples tend to accumulate where the gradient of the function is high.

3. Restricted Voronoi diagrams

Given a set of samples $\mathbf{V} = \{\mathbf{v}_k\}_k \subset \mathbb{R}^3$ and a surface $\mathcal{S} \subset \mathbb{R}^3$, a restricted Voronoi cell $\Omega_{k|\mathcal{S}}$ is defined as:

$$\Omega_{k|\mathcal{S}} = \{\mathbf{x} \in \mathcal{S}, \|\mathbf{x} - \mathbf{v}_k\| < \|\mathbf{x} - \mathbf{v}_\ell\| \text{ for all } \mathbf{v}_\ell \in \mathbf{V}\}. \quad (1)$$

In other words, it corresponds to the intersection between the 3d Voronoi cell of a sample \mathbf{v}_k and the mesh \mathcal{S} . The restricted Voronoi cells are polygonal meshes. They define a partition of \mathcal{S} . The set of all the restricted Voronoi cells and their boundaries is called the restricted Voronoi diagram. An example is shown on Figure 2.

The set of samples \mathbf{V} is said to be in general position if it

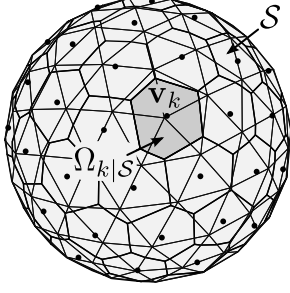


Figure 2: *Restricted Voronoi Diagram.* The restricted Voronoi cell $\Omega_{k|\mathcal{S}}$ is the set of points $\mathbf{x} \in \mathcal{S}$ for which \mathbf{v}_k is the nearest sample.

generically defines a Voronoi diagram. In other words, three samples are not colinear, four samples are not coplanar, four samples are not on a same circle, and five samples are not on a same sphere. In addition, we require that the intersection between the vertices, edges and faces of \mathcal{S} , and the vertices, edges and faces of the Voronoi cells are either empty or have the right dimension. For instance, no vertex of \mathcal{S} is contained in the boundary of a Voronoi cell, and the intersection between a Voronoi face and an edge of \mathcal{S} is either a point or empty. In practice, symbolic perturbation [EM90] can be used to handle degenerate configurations.

4. Objective function

4.1. Building an approximation

Our goal in this work is to compute and approximation of a function $f : \mathcal{S} \rightarrow \mathbb{R}$ defined on a mesh \mathcal{S} . We here assume that f is bounded, piecewise continuous and locally integrable. Our approximation is encoded by

- a set of n sample positions $\mathbf{V} = \{\mathbf{v}_k\}_k \subset \mathbb{R}^3$;
- a set of n sample approximation values $\tilde{\mathbf{F}} = \{\tilde{f}_k\}_k \subset \mathbb{R}$.

From this data, our approximation $\tilde{f} : \mathcal{S} \rightarrow \mathbb{R}$ is defined by mapping each point $\mathbf{x} \in \mathcal{S}$ to the approximation value \tilde{f}_k of its nearest sample \mathbf{v}_k . The main advantage of this method is that the data stored for the approximation is independent from the combinatorics of \mathcal{S} . In particular, it decouples the sampling resolution from the mesh resolution, which is suitable for applications like light sampling [LZT*08]. The approximation \tilde{f} is however piecewise constant on \mathcal{S} , and therefore discontinuous.

From \mathbf{V} and $\tilde{\mathbf{F}}$, the value at a given point $\mathbf{x} \in \mathcal{S}$ can be retrieved using a nearest neighbor search, which can be done efficiently using space decomposition techniques [MA97]. The approximation can also be computed globally by computing the Voronoi diagram of \mathbf{V} restricted to \mathcal{S} . Each region $\Omega_{k|\mathcal{S}}$ of \mathcal{S} (see Equation 1) indeed corresponds to the portion of \mathcal{S} that receives the approximation value \tilde{f}_k . Such a restricted Voronoi diagram can be efficiently and exactly computed using the algorithm of Yan et al. [YLL*09], or approximately using GPU techniques [HIKL*99, SGG*07].

4.2. Objective function

To optimize the placement of the samples, we minimize the objective function defined below. The rationale for using this specific function is explained further.

$$\mathcal{F}(\mathbf{V}) = - \sum_{k=1}^n \frac{1}{|\Omega_{k|\mathcal{S}}|} \left(\int_{\Omega_{k|\mathcal{S}}} f(\mathbf{x}) d\mathbf{x} \right)^2, \quad (2)$$

where $|\Omega_{k|\mathcal{S}}|$ is the area of the restricted Voronoi cell $\Omega_{k|\mathcal{S}}$. Note that this objective function solely depends on the sample positions \mathbf{V} . The approximation values $\tilde{\mathbf{F}}$ are then computed from the sample positions as :

$$\tilde{f}_k = \frac{1}{|\Omega_{k|\mathcal{S}}|} \int_{\Omega_{k|\mathcal{S}}} f(\mathbf{x}) d\mathbf{x}. \quad (3)$$

We now further explain the relation between this objective function and the initial problem. The quality of an approximation can be measured by summing the squared difference between f and \tilde{f} at each point $\mathbf{x} \in \mathcal{S}$:

$$\mathcal{G}(\mathbf{V}, \tilde{\mathbf{F}}) = \int_{\mathcal{S}} \|f(\mathbf{x}) - \tilde{f}(\mathbf{x})\|^2 d\mathbf{x}.$$

Segmenting \mathcal{S} using the Voronoi diagram of \mathbf{V} , this function can be rewritten as :

$$\mathcal{G}(\mathbf{V}, \tilde{\mathbf{F}}) = \sum_{k=1}^n \int_{\Omega_{k|\mathcal{S}}} \|f(\mathbf{x}) - \tilde{f}_k\|^2 d\mathbf{x}. \quad (4)$$

Considering the sample positions \mathbf{V} as constants, \mathcal{G} falls back to a classical least squares minimization problem, which is minimized when each approximation value \tilde{f}_k coincides with the mean value of f on $\Omega_{k|\mathcal{S}}$. From now on, we will therefore consider that the approximation values are bound to the sample positions \mathbf{V} by Equation 3. With this assumption, our minimization problem reduces to finding a set of samples \mathbf{V} that minimize the objective function :

$$\mathcal{G}(\mathbf{V}) = \sum_{k=1}^n \int_{\Omega_{k|\mathcal{S}}} \|f(\mathbf{x}) - \tilde{f}_k(\mathbf{V})\|^2 d\mathbf{x},$$

which corresponds to the *variance* of f on the restricted Voronoi cells. The function \mathcal{G} can be reformulated as :

$$\begin{aligned} \mathcal{G}(\mathbf{V}) &= \sum_{k=1}^n \left[\int_{\Omega_{k|\mathcal{S}}} f(\mathbf{x})^2 d\mathbf{x} - \tilde{f}_k(\mathbf{V})^2 |\Omega_{k|\mathcal{S}}| \right] \\ &= \int_{\mathcal{S}} f(\mathbf{x})^2 d\mathbf{x} - \sum_{k=1}^n \frac{1}{|\Omega_{k|\mathcal{S}}|} \left(\int_{\Omega_{k|\mathcal{S}}} f(\mathbf{x}) d\mathbf{x} \right)^2 \\ &= \int_{\mathcal{S}} f(\mathbf{x})^2 d\mathbf{x} + \mathcal{F}(\mathbf{V}). \end{aligned} \quad (5)$$

Finally, since the first term does not depend on the sample positions, it can be discarded without changing the minimizer of \mathcal{G} , hence the final expression of our objective function, as stated by Equation 2.

A *global* minimum of \mathcal{F} is very difficult to obtain. We therefore focus on finding a *local* minimum of \mathcal{F} . To do so,

we will use a variant of a *gradient descent* method, described in Section 5. Such a method needs to evaluate the gradient of \mathcal{F} , we will therefore now focus on its computation. In other words, we need to compute gradients of expressions with *variable integration domains*. A key mathematical tool for that is Reynolds transport theorem presented in the next section.

4.3. Reynolds transport theorem

We quickly review existing methods that compute the gradient of \mathcal{F} with respect to the sample positions. The difficulty lies in the fact that the integration domains $\Omega_k|_{\mathcal{S}}$ vary with the sample positions. Lévy and Liu [LL10] decompose $\Omega_k|_{\mathcal{S}}$ into triangles, are able to provide an analytic formula for the integral. The gradient is then derived from this formula. In our context, such an analytic formula is not available since we suppose that f can only be evaluated at discrete locations. Carrier Baudouin et al. [CBRM*12] use a parameterization $\mathbf{x}(u, v)$ of each triangle onto a reference triangle which no longer depends on \mathbf{V} . The derivation with respect to a sample can therefore be transferred under the integrals. The problem however is that the derivation of $f(\mathbf{x}(u, v))$ requires the value of the derivative of f at any point $\mathbf{x} \in \mathcal{S}$.

For this reason, we use another technique to compute the derivative of \mathcal{F} , with *Reynolds transport theorem* [Rey03]. A formulation of this theorem was proved by Cortes, Martinez and Bullo [CMB05] in the context of the optimization of 2d Voronoi cells, and [dGBOD12] also use this theorem to compute the derivative of their objective function. If F is a face of \mathcal{S} , and $\Omega_k|_F$ the Voronoi cell of \mathbf{v}_k restricted to F , and if $\phi : \mathcal{S} \times \mathbb{R}^{3n} \rightarrow \mathbb{R}$ is a function of a point $\mathbf{x} \in \mathcal{S}$ and the set of samples \mathbf{V} , continuous on $\mathcal{S} \times \mathbb{R}^{3n}$, continuously differentiable with respect to the sample positions, and provided the following integrals are well defined, then :

$$\frac{d}{d\mathbf{v}_\ell} \left[\int_{\Omega_k|_F} \phi(\mathbf{x}, \mathbf{V}) d\mathbf{x} \right] = \int_{\Omega_k|_F} \frac{\partial \phi}{\partial \mathbf{v}_\ell}(\mathbf{x}, \mathbf{V}) d\mathbf{x} + \int_{\partial \Omega_k|_F} \phi(\mathbf{x}, \mathbf{V}) \frac{d\mathbf{x}}{d\mathbf{v}_\ell} \cdot \mathbf{n}_b d\mathbf{x}, \quad (6)$$

where $\partial \Omega_k|_F$ is the boundary of $\Omega_k|_F$, $\frac{d\mathbf{x}}{d\mathbf{v}_\ell}$ encodes the variations of this boundary with respect to the position of the sample \mathbf{v}_ℓ , and \mathbf{n}_b is the normal of $\partial \Omega_k|_F$ at \mathbf{x} . Appendix B extends this theorem when f is only piecewise smooth on \mathcal{S} . We will first focus on the computation of $\frac{d\mathbf{x}}{d\mathbf{v}_\ell} \cdot \mathbf{n}_b$.

4.4. Boundary variations

Let us assume the samples are in a general position \mathbf{V} . Let $[\mathbf{c}_1, \mathbf{c}_2]$ be an edge of the polygon corresponding to $\Omega_k|_F$, this section details the following result :

- if $[\mathbf{c}_1, \mathbf{c}_2]$ is the intersection between an edge of F and Ω_k

then for any point $\mathbf{x} \in [\mathbf{c}_1, \mathbf{c}_2]$ and any sample \mathbf{v}_ℓ

$$\frac{d\mathbf{x}}{d\mathbf{v}_\ell} \cdot \mathbf{n}_b = \mathbf{0}^t \quad (7)$$

- otherwise, $[\mathbf{c}_1, \mathbf{c}_2]$ is the intersection between F and a Voronoi face separating Ω_k and another cell Ω_ℓ :

$$\frac{d\mathbf{x}}{d\mathbf{v}_k} \cdot \mathbf{n}_b = \frac{(\mathbf{v}_k - \mathbf{x})^t}{\mathbf{n}_b \cdot (\mathbf{v}_k - \mathbf{v}_\ell)}, \quad \frac{d\mathbf{x}}{d\mathbf{v}_\ell} \cdot \mathbf{n}_b = \frac{(\mathbf{x} - \mathbf{v}_\ell)^t}{\mathbf{n}_b \cdot (\mathbf{v}_k - \mathbf{v}_\ell)}, \quad (8)$$

and the derivative with respect to any other sample is zero.

The fact that the samples are in general position means that there exists a neighborhood N of \mathbf{V} in which the combinatorics of the restricted Voronoi diagram does not change. Let $\mathbf{x} : [0, 1] \times N \rightarrow \mathcal{S}$ be a parametrization of the edge $[\mathbf{c}_1, \mathbf{c}_2]$ on the unit segment $[0, 1]$ for any sample configuration $\mathbf{V} \in N$. Let $d\mathbf{V} = \{d\mathbf{v}_k\}_k \in \mathbb{R}^{3n}$ be a non zero perturbation of the sample positions and let $h \in \mathbb{R}^+$. The derivative of a point \mathbf{x} of $[\mathbf{c}_1, \mathbf{c}_2]$ in the direction $d\mathbf{V}$ is given by the limit

$$\frac{d\mathbf{x}}{d\mathbf{V}}(u, \mathbf{V}) = \lim_{h \rightarrow 0} \frac{\mathbf{x}(u, \mathbf{V} + h d\mathbf{V}) - \mathbf{x}(u, \mathbf{V})}{h}. \quad (9)$$

Without loss of generality, we consider h to be small enough so that $\mathbf{V} + h d\mathbf{V} \in N$. Since the dot product bilinear form is clearly continuous in its own arguments, we have

$$\frac{d\mathbf{x}}{d\mathbf{V}}(u, \mathbf{V}) \cdot \mathbf{n}_b = \lim_{h \rightarrow 0} \left[\frac{\mathbf{x}(u, \mathbf{V} + h d\mathbf{V}) - \mathbf{x}(u, \mathbf{V})}{h} \cdot \mathbf{n}_b \right]. \quad (10)$$

The edge $[\mathbf{c}_1, \mathbf{c}_2]$ can be either the intersection of an edge of F and Ω_k or the intersection between F and a face of Ω_k . In the first case, the points $\mathbf{x}(u, \mathbf{V})$ and $\mathbf{x}(u, \mathbf{V} + h d\mathbf{V})$ both belong to the edge of F , and the difference of these two points is therefore a vector colinear with the edge, and orthogonal to the edge normal, hence the result provided by Equation 7. If $[\mathbf{c}_1, \mathbf{c}_2]$ is the intersection between a Voronoi face and F , the plane of the Voronoi face is the bisector between \mathbf{v}_k and another sample \mathbf{v}_ℓ . Let a be a scalar and \mathbf{w} be any unit vector such that :

$$\mathbf{x}(u, \mathbf{V} + h d\mathbf{V}) - \mathbf{x}(u, \mathbf{V}) = a \mathbf{w}. \quad (11)$$

Since $\mathbf{x}(u, \mathbf{V} + h d\mathbf{V})$ is contained by the bisector of $\mathbf{v}_k + h d\mathbf{v}_k$ and $\mathbf{v}_\ell + h d\mathbf{v}_\ell$, and since $\mathbf{x}(u, \mathbf{V})$ is in the bisector of \mathbf{v}_k and \mathbf{v}_ℓ , we have :

$$a = \frac{h((\mathbf{v}_k - \mathbf{x}(u, \mathbf{V})) \cdot d\mathbf{v}_k - (\mathbf{v}_\ell - \mathbf{x}(u, \mathbf{V})) \cdot d\mathbf{v}_\ell) + o(h)}{\mathbf{w} \cdot (\mathbf{v}_k - \mathbf{v}_\ell) + O(h)}. \quad (12)$$

If \mathbf{e} is a unit vector contained in $[\mathbf{c}_1, \mathbf{c}_2]$, and \mathbf{n}_b is the normal of $[\mathbf{c}_1, \mathbf{c}_2]$ in F , the vectors \mathbf{e} and \mathbf{n}_b form an orthogonal basis of the plane of F , and we can write :

$$\mathbf{w} = (\mathbf{w} \cdot \mathbf{e}) \mathbf{e} + (\mathbf{w} \cdot \mathbf{n}_b) \mathbf{n}_b.$$

In addition, $[\mathbf{c}_1, \mathbf{c}_2]$ is contained in the bisector of \mathbf{v}_k and \mathbf{v}_ℓ , thus $\mathbf{e} \cdot (\mathbf{v}_k - \mathbf{v}_\ell) = 0$. Therefore,

$$a \mathbf{w} \cdot \mathbf{n}_b = \frac{h((\mathbf{v}_k - \mathbf{x}(u, \mathbf{V})) \cdot d\mathbf{v}_k - (\mathbf{v}_\ell - \mathbf{x}(u, \mathbf{V})) \cdot d\mathbf{v}_\ell) + o(h)}{\mathbf{n}_b \cdot (\mathbf{v}_k - \mathbf{v}_\ell) + O(h)}. \quad (13)$$

Note that this expression of $aw \cdot \mathbf{n}_b$ no longer depends on the direction \mathbf{w} . This remark will be useful to study the continuity of \mathcal{F} (see Equation 19, Appendix A). By substituting the expression of $a\mathbf{w}$ in Equation 10 and Equation 11, we obtain the expression of the gradient (Equation 8). A property worth a remark is that this result does not depend on the parametrization chosen for $[\mathbf{c}_1, \mathbf{c}_2]$ as well.

4.5. Gradient of \mathcal{F}

With the variations of the boundary derived in Equations 7 and 8, we are now able to express the derivative of our objective function with respect to a given sample \mathbf{v}_ℓ . Supposing that f is evaluated at \mathbf{x} , and that $\frac{d\mathcal{F}}{d\mathbf{v}_\ell}$, \tilde{f}_k and \tilde{f}_ℓ are evaluated at a position \mathbf{V} of the samples,

$$\frac{d\mathcal{F}}{d\mathbf{v}_\ell} = \sum_{p,q} \int_{[\mathbf{c}_p, \mathbf{c}_q]} (\|f - \tilde{f}_\ell\|^2 - \|f - \tilde{f}_k\|^2) \frac{(\mathbf{v}_\ell - \mathbf{x})^t}{\mathbf{n}_b^\ell \cdot (\mathbf{v}_\ell - \mathbf{v}_k)} d\mathbf{x}, \quad (14)$$

where $[\mathbf{c}_p, \mathbf{c}_q]$ is an edge on the boundary of $\Omega_{\ell|\mathcal{S}}$, \mathbf{v}_k is the sample neighboring \mathbf{v}_ℓ along this edge, and \mathbf{n}_b^ℓ is the normal of the edge oriented outside of $\Omega_{\ell|\mathcal{S}}$.

Proof This expression is obtained from Equation 2 by expanding the derivatives, which leads to :

$$\frac{d\mathcal{F}}{d\mathbf{v}_\ell} = \sum_k \left[\tilde{f}_k^2 \frac{d|\Omega_{k|\mathcal{S}}|}{d\mathbf{v}_\ell} - 2\tilde{f}_k \frac{d}{d\mathbf{v}_\ell} \left(\int_{\Omega_{k|\mathcal{S}}} f(\mathbf{x}) d\mathbf{x} \right) \right] \quad (15)$$

Using Reynolds transport theorem with $\phi(\mathbf{x}, \mathbf{V}) = 1$ and $\phi(\mathbf{x}, \mathbf{V}) = f(\mathbf{x})$ respectively, in both cases the first term of Equation 6 vanishes, since ϕ does not depend on \mathbf{V} . Among all the edges of the restricted Voronoï diagram, the only edges that depend on \mathbf{v}_ℓ are those on the boundary of $\Omega_{\ell|\mathcal{S}}$. If $[\mathbf{c}_p, \mathbf{c}_q]$ is such an edge between restricted cells $\Omega_{k|\mathcal{S}}$ and $\Omega_{\ell|\mathcal{S}}$, the term on $\Omega_{k|\mathcal{S}}$ in the sum of Equation 15 yields the gradient :

$$\int_{[\mathbf{c}_p, \mathbf{c}_q]} (\|f - \tilde{f}_k\|^2 - f^2) \frac{(\mathbf{x} - \mathbf{v}_\ell)^t}{\mathbf{n}_b^k \cdot (\mathbf{v}_k - \mathbf{v}_\ell)} d\mathbf{x},$$

where \mathbf{n}_b^k is the normal of $[\mathbf{c}_p, \mathbf{c}_q]$ oriented outside of $\Omega_{k|\mathcal{S}}$, and the term on $\Omega_{\ell|\mathcal{S}}$ gives us

$$\int_{[\mathbf{c}_p, \mathbf{c}_q]} (\|f - \tilde{f}_\ell\|^2 - f^2) \frac{(\mathbf{v}_\ell - \mathbf{x})^t}{\mathbf{n}_b^\ell \cdot (\mathbf{v}_\ell - \mathbf{v}_k)} d\mathbf{x}.$$

Grouping these two integrals on the same edge, and using the fact that $\mathbf{n}_b^\ell = -\mathbf{n}_b^k$, we obtain the result presented in Equation 14. \square

4.6. Continuity

We here summarize our results about the continuity of \mathcal{F} and delay the formal proof to Appendix A. Studying the continuity of \mathcal{F} is important since it determines the class of solvers which can be used to minimize it. We here suppose that f is continuous.

C^0 continuity : \mathcal{F} is C^0 continuous almost everywhere. The only situations when it is not are when a face of \mathcal{S} and a Voronoï face of the diagram of \mathbf{V} intersect with dimension 2. This is proved as Theorem 1 in Appendix A.

C^1 continuity : \mathcal{F} is C^1 continuous almost whenever it is C^0 . The additional discontinuities happen when an edge of \mathcal{S} intersects the Voronoï faces with dimension 1, or when a Voronoï edge intersects a face of \mathcal{S} with dimension 1. This is proved as Theorem 2 in Appendix A.

These results render our objective function eligible for gradient descent methods, but not for higher order methods based on the Hessian of the objective function.

5. Solution mechanism

5.1. Gradient descent

The gradient descent algorithm that we use finds a local minimum of an objective function by successively updating the positions of the samples [NW99] :

$$\mathbf{V}^{(i+1)} = \mathbf{V}^{(i)} - \delta^{(i)} \left(\frac{d\mathcal{F}}{d\mathbf{V}}(\mathbf{V}^{(i)}) \right)^t,$$

where i is the index of the current step and $\delta^{(i)}$ is a positive scalar called the *step length*, which controls the amplitude of the update of the sample positions. To determine an adequate step length value, a common method is to use a *linesearch* algorithm, which ensures that each step decreases the value of the objective function. Using such a technique however does not work well in our case, since the objective function can be very anisotropic, in the sense that there is a wide variation of the magnitudes of the derivatives with respect to different variables. In addition, when f has discontinuities, a minimum may correspond to a situation when the gradient is discontinuous, which causes the linesearch to fail in fulfilling usual conditions like Wolfe conditions [NW99]. For this reason, we use instead the following heuristic to compute the step length at each iteration step :

$$\delta^{(i)} = \delta_0 \sigma^{r(i)}, \quad \text{with} \quad r(i) = \frac{i}{i_{max} - i}, \quad (16)$$

where i_{max} is the number of iterations required by the user, δ_0 is the initial step length, and σ controls how fast the step length decreases to 0. In all of the results presented in Section 6, we used the same values for these parameters :

- δ_0 is 2% of the scene bounding box diagonal ;
- $\sigma = \frac{1}{2}$.

In addition, to limit the effects of the anisotropy of the objective function, we do not use the global gradient direction for the update, but use instead the *normalized gradient direction* for each sample :

$$\mathbf{v}_k^{(i+1)} = \mathbf{v}_k^{(i)} + \delta^{(i)} \frac{\mathbf{g}_k^{(i)}}{\|\mathbf{g}_k^{(i)}\|}, \quad \text{with} \quad \mathbf{g}_k^{(i)} = \left(\frac{d\mathcal{F}}{d\mathbf{v}_k}(\mathbf{V}^{(i)}) \right)^t \quad (17)$$

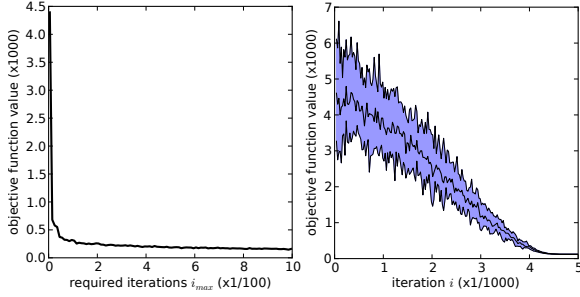


Figure 3: Convergence of the solver. These curves were generated using the function presented in Figure 5 with 1000 samples. Left, the value of the objective function reached as a function of the number of iterations i_{max} requested by the user. Right, the evolution of the objective function value with the current iteration in a single run of 5000 iterations. As no linesearch is performed, the objective function oscillates within the blue area, and the center curve is the median of the objective function values over periods of 25 successive iterations.

Figure 3 illustrates the convergence of this iteration scheme. Since no linesearch mechanism is used, the objective function does not necessarily decrease at each iteration and oscillates. The oscillations decrease with the step length.

5.2. Approximation of the integrals

We consider f as a *black box*. By black box, we mean that no closed form formula may be available for f . The only requirement we have for f is a way of evaluating it at any point of \mathcal{S} . To estimate the integrals, quadrature rules are needed [PTVF07]. Integrals are required at two different steps. First, we need to compute the approximation values of the samples at each step, using Equation 3. Then, to compute the gradient of \mathcal{F} given in Equation 14, we need integrals on the edges of the restricted Voronoi diagram.

In our experiments, quadratures of degree two have proved sufficient. Each restricted Voronoi cell is decomposed into triangles, and f is evaluated at three points per triangle [Dun85]. For the gradients, f is evaluated at two points per edge of the restricted Voronoi diagram.

5.3. Algorithm

Our numerical solution mechanism is summarized in Algorithm 1. In terms of complexity, each step requires the computation of the Voronoi diagram of the samples and the computation of the restricted Voronoi cells. Although no theoretical complexity analysis is available, the empirical results presented by Yan, Lévy et al. [YLL*09] tend to show that this algorithm approximately behaves with a $O(m+n)$ complexity in most practical scenarios. Figure 4 confirms this behavior in our case.

input : a mesh \mathcal{S} , a function f on \mathcal{S} , the number of samples n and iterations i_{max} .
output: optimized samples and approximation values.

```

Generate  $n$  random samples on  $\mathcal{S}$  for  $\mathbf{V}^{(0)}$ 
foreach  $i$  from 0 to  $i_{max}$  do
    Compute restricted Voronoi diagram of  $\mathbf{V}^{(i)}$  and  $\mathcal{S}$ 
    foreach restricted cell  $\Omega_k | \mathcal{S}$  do
        // Compute the sample values with Equation 3
         $\tilde{f}_k \leftarrow \frac{1}{|\Omega_k | \mathcal{S}|} \int_{\Omega_k | \mathcal{S}} f(\mathbf{x}) d\mathbf{x}$ 
    foreach edge of  $\Omega_k | \mathcal{S} \cap \Omega_{\ell} | \mathcal{S}$ , with  $k < \ell$  do
        // Accumulate the gradient using Equation 14
        update  $\frac{d\mathcal{F}}{d\mathbf{v}_k}$  and  $\frac{d\mathcal{F}}{d\mathbf{v}_\ell}$ 
    Compute the step length  $\delta^{(i)}$  using Equation 16
    foreach sample  $\mathbf{v}_k$  do
        // Update the sample positions
         $\mathbf{v}_k^{(i+1)} = \mathbf{v}_k^{(i)} + \delta^{(i)} g_k^{(i)}$ 
    
```

Algorithm 1: Optimization of a set of samples to approximate a function f on a mesh \mathcal{S} .

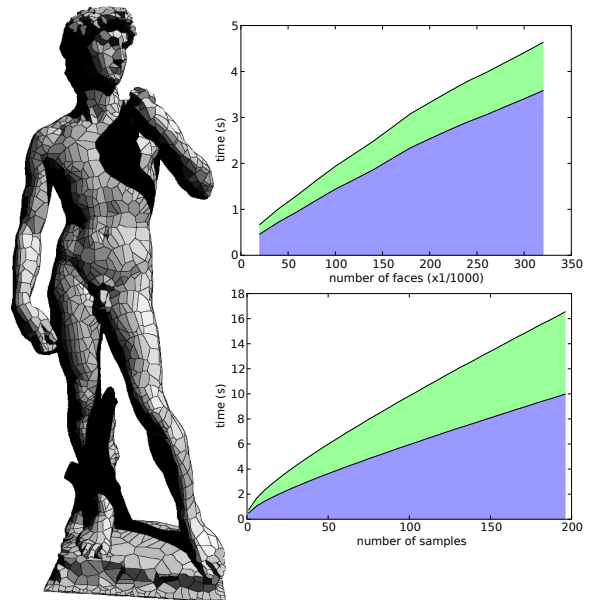


Figure 4: Timing of one gradient descent algorithm step with respect to the resolution of the mesh and the number of samples. The approximated function is the direct lighting of the David mesh, and the various resolutions were obtained via a remeshing algorithm. The blue regions correspond to the portion of the time dedicated to the computation of the approximation values with triangle integrals. The green part is the time dedicated to the gradient computation, with line integrals.




Figure 5: Approximation of the function $\arctan\left(\frac{y}{x}\right)$ on $[-1, 1]^2$. The top shows the approximation, the diagram is displayed at the bottom, and a close-up of the center is provided on the right. The samples align along the direction of the gradient of f and on both sides of the discontinuity, to ensure that the boundary of the Voronoi cells match the discontinuity. 500 samples, 1500 iterations for 20 sec runtime

6. Results

6.1. Test function

We first demonstrate the behavior of our algorithm on a simple analytic function (Figure 5). The function we use is $f(\mathbf{x}) = \arctan\left(\frac{y}{x}\right)$, which is piecewise continuous except for $x = y = 0$ and exhibits a discontinuity along the $x = 0$ line. The result obtained by the minimization exhibits the following characteristic properties :

the samples align along gradient directions. as a consequence, the cells are thin in the direction where f varies most ;

pairs of sample capture the discontinuities by aligning on each side of the discontinuity, in such a way that the edge between the cells is along the discontinuity.

As explained in Appendix B, when f has straight discontinuities, the minimum of the objective function is a point where the gradient of \mathcal{F} is discontinuous, since edges of the restricted Voronoi diagram align with the discontinuities. This also explains why usual linesearch techniques fail to efficiently minimize our objective function.

6.2. Image sampling

Figure 6 demonstrates our method used in the context of image vectorization. In this context, the function f is defined by an rgb image and is thus trivariate. It is used as a black box in the sense that it is only queried through the evaluations required by the quadrature rules mentioned in Section 5.2. In particular, we emphasize that Voronoi cell boundaries are naturally aligned with the discontinuities. This behavior naturally emerges from optimizing our objective function, without requiring any edge detection phase. As for Figure 5, the

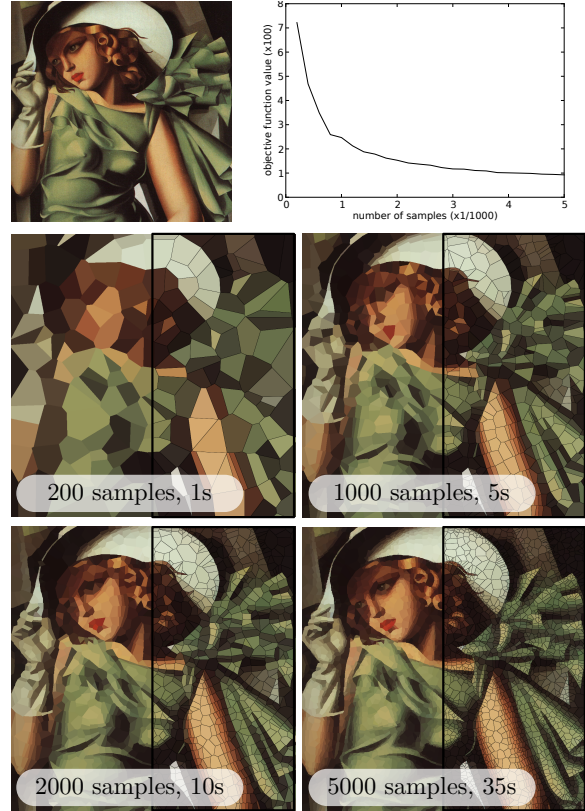


Figure 6: Vectorizing a color painting by Tamara de Lempicka (top left). Sampling size and timings are given on each result, as well as a study of the final objective function value with respect to the number of samples. All the results were obtained using 300 iterations.



Figure 7: Comparison between a random sampling (left) of an image and our result (right), both using 1000 samples. Objective function values are 0.09 for the random sampling and 0.02 for the optimized one.

sharp features are well captured. Using the parameters given in Section 5.1, the algorithm is fully automatic. We also plot in Figure 6 the evolution of the final value of \mathcal{F} obtained with respect to the number of samples optimized using 300 iterations. Figure 7 demonstrates the difference between a random sampling and our result on a greylevel image.




Figure 8: Direct illumination from a point source on the Pegasus model. The shadow boundaries are well captured, and very few samples are placed in regions with uniform color. 1500 samples, 300 iterations, 12 min runtime for a mesh with 85k faces.

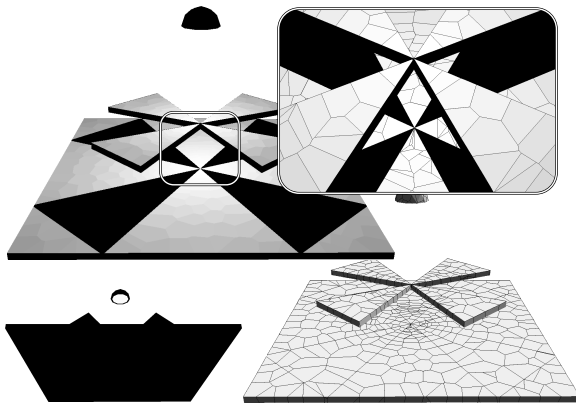


Figure 9: Direct lighting on a thin surface. The samples have to align on both sides of the surface to prevent the restricted Voronoi cells from having several components. Shadows remain well captured, even at the sharp angles in the center. 1500 samples, 500 iterations, 1 min runtime.

6.3. Direct illumination sampling

We now switch to 3d and show some examples of function approximation on 3d meshes. The function we use is the result of the direct illumination of the surfaces from a point source, with ray traced shadows. Figure 8 shows such an example. An additional difficulty on meshes appears because of the use of restricted Voronoi cells. When the mesh has thin features, the restricted Voronoi cells may have several connected components. In the context of lighting, such a phenomenon is particularly undesirable since the connected components may have opposite normals, one component being lit while the other remains in the shadow. As can be seen in Figure 9, our method properly handles these difficult configurations. The optimization of our objective function naturally places a sample on each side of the thin feature, which allows properly representing the lighting variation on them.

7. Conclusion and future work

We proposed a fully automatic method to approximate a function on a mesh. This method has very few requirements on the function to be approximated, since only evaluation is required. It makes it possible to use the method for functions that are not known explicitly, such as direct lighting in Figure 9, where evaluation is the result of an algorithm (e.g. ray-traced shadows). In its current state, the solver we use to optimize the sample positions could probably be improved to accelerate the convergence. Designing a new solver may however be difficult, since the optimal solution to the problem is a situation where the gradient of the objective function is discontinuous. Another limitation of our method is that the approximation we compute is only piecewise constant. We therefore plan on extending this approach to be able to generate higher order smooth approximations. An idea for such an extension could be done by using a method similar to the work of Orzan et al. [OBW*08], by specifying constraints for the diffusion process from the sample values and positions, and using a threshold on the edges of the restricted Voronoi diagram to create curves blocking the diffusion along discontinuities.

Acknowledgments

The authors wish to thank Nicolas Ray and Laurent Alonso for many discussions, proof reading and support. Bruno Levy is also partly supported by the European Research Council (GOODSHAPE ERC-StG-205693) and the ANR (MORPHO and BECASIM).

References

- [BB*12] BRUNO L., BONNEEL N., ET AL.: Variational anisotropic surface meshing with voronoi parallel linear enumeration. In *IMR-21st International Meshing Roundtable-2012* (2012). 2
- [CBRM*12] CARRIER BAUDOUIN T., REMACLE J.-F. C., MARCHANDISE E., LAMBRECHTS J., HENROTTE F.: Lloyd's energy minimization in the L_p norm for quadrilateral surface

- mesh generation. *Engineering with Computers* (2012), 1–14. doi:10.1007/s00366-012-0290-x. 4
- [CMB05] CORTES J., MARTINEZ S., BULLO F.: Spatially-distributed coverage optimization and control with limited-range interactions. *ESAIM contrôle optimisation et calcul des variations 11* (2005), 691. 4, 10
- [CSAD04] COHEN-STEINER D., ALLIEZ P., DESBRUN M.: Variational shape approximation. In *ACM Transactions on Graphics (TOG)* (2004), vol. 23, ACM, pp. 905–914. 2
- [dGBOD12] DE GOES F., BREEDEN K., OSTROMOUKHOV V., DESBRUN M.: Blue noise through optimal transport. *ACM Transactions on Graphics (TOG)* 31, 6 (2012), 171. 2, 4
- [Dun85] DUNAVANT D.: High degree efficient symmetrical gaussian quadrature rules for the triangle. *International journal for numerical methods in engineering* 21, 6 (1985), 1129–1148. 6
- [EM90] EDELSBRUNNER H., MÜCKE E. P.: Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM TRANS. GRAPH* 9, 1 (1990), 66–104. 3
- [FH04] FELZENSZWALB P. F., HUTTENLOCHER D. P.: Efficient graph-based image segmentation. *International Journal of Computer Vision* 59, 2 (2004), 167–181. 2
- [HA*04] HOLZSCHUCH N., ALONSO L., ET AL.: Combining higher-order wavelets and discontinuity meshing: A compact representation for radiosity. In *Eurographics Symposium on Rendering* (2004), pp. 275–286. 2
- [Hec92] HECKBERT P.: Discontinuity meshing for radiosity. In *Third Eurographics Workshop on Rendering* (1992), pp. 203–226. 2
- [HEH05] HOIEM D., EFROS A. A., HEBERT M.: Automatic photo pop-up. In *ACM Transactions on Graphics (TOG)* (2005), vol. 24, ACM, pp. 577–584. 2
- [HIKL*99] HOFF III K. E., KEYSER J., LIN M., MANOCHA D., CULVER T.: Fast computation of generalized voronoi diagrams using graphics hardware. In *Proceedings of the 26th annual conference on Computer graphics and interactive techniques* (1999), ACM Press/Addison-Wesley Publishing Co., pp. 277–286. 3
- [LL06] LECOT G., LÉVY B.: Ardeco: Automatic region detection and conversion. In *Proceedings of the 17th Eurographics conference on Rendering Techniques* (2006), Eurographics Association, pp. 349–360. 2
- [LL10] LÉVY B., LIU Y.: Lp centroidal voronoi tessellation and its applications. *ACM Transactions on Graphics* 29 (2010), 4. 4
- [LTG92] LISCHINSKI D., TAMPRIERI F., GREENBERG D. P.: Discontinuity meshing for accurate radiosity. *IEEE Computer graphics and Applications* 12, 6 (1992), 25–39. 2
- [LWL*09] LIU Y., WANG W., LÉVY B., SUN F., YAN D., LU L., YANG C.: On centroidal voronoi tessellation – energy smoothness and fast computation. *ACM Transactions on Graphics (ToG)* 28, 4 (2009), 101. 2
- [LZT*08] LEHTINEN J., ZWICKER M., TURQUIN E., KONTKANEN J., DURAND F., SILLION F. X., AILA T.: A meshless hierarchical representation for light transport. In *ACM Transactions on Graphics (TOG)* (2008), vol. 27, ACM, p. 37. 2, 3
- [MA97] MOUNT D. M., ARYA S.: ANN: A library for approximate nearest neighbor searching. In *CG Workshop on Computational Geometry* (1997), pp. 33–40. 3
- [NW99] NOCEDAL J., WRIGHT S.: *Numerical optimization*. Springer verlag, 1999. 5
- [OBW*08] ORZAN A., BOUSSEAU A., WINNEMÖLLER H., BARLA P., THOLLOT J., SALESIN D.: Diffusion curves: a vector representation for smooth-shaded images. In *ACM Transactions on Graphics (TOG)* (2008), vol. 27, ACM, p. 92. 2, 8
- [OS97] OKABE A., SUZUKI A.: Locational optimization problems solved through voronoi diagrams. *European Journal of Operational Research* 98, 3 (1997), 445–456. 2
- [PTVF07] PRESS W. H., TEUKOLSKY S. A., VETTERLING W. T., FLANNERY B. P.: *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press, 2007. 6
- [PV93] POCCHIOLA M., VEGTER G.: The visibility complex. In *Proceedings of the ninth annual symposium on Computational geometry* (1993), ACM, pp. 328–337. 2
- [Rey03] REYNOLDS O.: *Papers on Mechanical and Physical Subjects*, vol. 3. The University Press, 1903. 4
- [SGG*07] SUD A., GOVINDARAJU N., GAYLE R., ANDERSEN E., MANOCHA D.: Surface distance maps. In *Proceedings of Graphics Interface 2007* (2007), ACM, pp. 35–42. 3
- [SLWS07] SUN J., LIANG L., WEN F., SHUM H.-Y.: Image vectorization using optimized gradient meshes. In *ACM Transactions on Graphics (TOG)* (2007), vol. 26, ACM, p. 11. 2
- [YLL*09] YAN D.-M., LÉVY B., LIU Y., SUN F., WANG W.: Isotropic Remeshing with Fast and Exact Computation of Restricted Voronoi Diagram. *Computer Graphics Forum* 28, 5 (July 2009), 1445–1454. doi:10.1111/j.1467-8659.2009.01521.x. 3, 6

Appendix A: Detailed continuity of \mathcal{F}

Theorem 1 If $\phi : \mathcal{S} \times \mathbb{R}^{3n} \rightarrow \mathbb{R}$ is a continuous function, and if $\Omega_{k|F}$ denotes the Voronoi cell of a sample \mathbf{v}_k restricted to a face F of \mathcal{S} , then

$$\mathcal{F}_{k|F}(\mathbf{V}) = \int_{\Omega_{k|F}} \phi(\mathbf{x}, \mathbf{V}) d\mathbf{x},$$

is continuous when F is not coplanar with a face of Ω_k .

Proof What we want to prove is that for any $\epsilon > 0$, there exists a $\rho > 0$ such that if the amplitude $h > 0$ of a displacement of the samples in any direction $d\mathbf{V}$ is such that $h < \rho$, we have

$$\|\Delta \mathcal{F}_{k|F}\| = \|\mathcal{F}_{k|F}(\mathbf{V} + h d\mathbf{V}) - \mathcal{F}_{k|F}(\mathbf{V})\| < \epsilon. \quad (18)$$

Whatever the combinatorics of the vertices of the boundary of $\Omega_{k|F}$, these vertices depend smoothly on the sample positions if F is not coplanar with a Voronoi face of \mathbf{v}_k . Therefore, for any $\epsilon_1 > 0$, there exists ρ_1 such that if $h < \rho_1$, the displacement of a vertex \mathbf{c}_1 of the boundary of $\Omega_{k|F}$ is contained in a ball of radius ϵ_1 . Let Λ_{ϵ_1} be the set of all points at distance at most ϵ_1 from $\partial\Omega_{k|F}$ and let $\Omega'_{k|F}$ denote the restricted Voronoi cell of $\mathbf{v}_k + h d\mathbf{v}_k$. We are therefore sure that no point of $\Omega'_{k|F}$ is outside of $\Omega_{k|F} \cup \Lambda_{\epsilon_1}$. We can thus write

$$\begin{aligned} \Delta \mathcal{F}_{k|F} &= \int_{\Omega'_{k|F} \setminus \Lambda_{\epsilon_1}} \phi(\mathbf{x}, \mathbf{V} + h d\mathbf{V}) d\mathbf{x} - \int_{\Omega_{k|F} \setminus \Lambda_{\epsilon_1}} \phi(\mathbf{x}, \mathbf{V}) d\mathbf{x} \\ &+ \int_{\Omega'_{k|F} \cap \Lambda_{\epsilon_1}} \phi(\mathbf{x}, \mathbf{V} + h d\mathbf{V}) d\mathbf{x} - \int_{\Omega_{k|F} \cap \Lambda_{\epsilon_1}} \phi(\mathbf{x}, \mathbf{V}) d\mathbf{x}. \end{aligned}$$

First, since ϕ is continuous on F , it is bounded on this domain by some value f_{max} , and the last two terms of this sum

can be bounded by

$$\int_{\Omega_{k|F} \cap \Lambda_{\varepsilon_1}} \phi(\mathbf{x}, \mathbf{V}) d\mathbf{x} \leq |\Lambda_{\varepsilon_1}| f_{max}$$

where $|\Lambda_{\varepsilon_1}|$ is the area of Λ_{ε_1} , and is $O(\varepsilon_1)$. In addition, the domains of the first two integrals are the same, since the boundary variations of $\Omega_{k|F}$ are enclosed in Λ_{ε_1} . These integrals can thus be combined to form

$$\int_{\Omega_{k|F} \setminus \Lambda_{\varepsilon_1}} \phi(\mathbf{x}, \mathbf{V} + h d\mathbf{V}) - \phi(\mathbf{x}, \mathbf{V}) d\mathbf{x}.$$

Since ϕ is continuous, for any $\varepsilon_2 > 0$, we can find a $\rho_2 > 0$ such that if $h < \rho_2$, $\phi(\mathbf{x}, \mathbf{V} + h d\mathbf{V}) - \phi(\mathbf{x}, \mathbf{V}) < \varepsilon_2$. We finally obtain

$$\Delta \mathcal{F}_{k|F} \leq |\Omega_{k|F}| \varepsilon_2 + 2|\Lambda_{\varepsilon_1}| f_{max} = O(\varepsilon_1 + \varepsilon_2).$$

Decreasing the values of ρ_1 and ρ_2 such that $\|\Delta \mathcal{F}_{k|F}\| < \varepsilon$, and using $\rho = \min(\rho_1, \rho_2)$, we obtain the result. \square

Theorem 2 If $\phi : \mathcal{S} \times \mathbb{R}^{3n} \rightarrow \mathbb{R}$ is a continuous function such that for each $\mathbf{x} \in \mathcal{S}$ the map $\phi_{\mathbf{x}} : \mathbf{V} \mapsto \phi(\mathbf{x}, \mathbf{V})$ is continuously differentiable, and if $\Omega_{k|F}$ denotes the Voronoi cell of a sample \mathbf{v}_k restricted to a face F of \mathcal{S} , then

$$\mathcal{F}_{k|F}(\mathbf{V}) = \int_{\Omega_{k|F}} \phi(\mathbf{x}, \mathbf{V}) d\mathbf{x},$$

is continuously differentiable if no edge of F is contained in a face of Ω_k and no edge of Ω_k is contained in F .

Proof We will prove that all the partial derivatives of \mathcal{F} exist and are continuous. Let us consider the partial derivative with respect to the sample \mathbf{v}_k . All the other samples are therefore considered as constants. Let us denote $\Gamma_{\ell|F}$ the restricted Voronoi cell of any other sample \mathbf{v}_ℓ in the Voronoi diagram of $\mathbf{V} \setminus \{\mathbf{v}_k\}$. We can thus write

$$\mathcal{F}_{k|F}(\mathbf{V}) = \sum_{\ell \neq k} \int_{\Gamma_{\ell|F} \cap \Omega_k} \phi(\mathbf{x}, \mathbf{V}) d\mathbf{x}.$$

Intersecting $\Gamma_{\ell|F}$ is actually just a matter of cutting it by the bisector of \mathbf{v}_k and \mathbf{v}_ℓ . If the bisector contains no vertex of $\Gamma_{\ell|F}$, a sufficiently small displacement of \mathbf{v}_k will not change the combinatorics of $\Gamma_{\ell|F} \cap \Omega_k$, and we can apply Reynolds transport theorem. The integral is therefore continuous in this situation.

If the bisector contains a vertex \mathbf{c}_1 , the locus of the samples \mathbf{v}_k such that \mathbf{c}_1 is contained in the bisector of \mathbf{v}_ℓ and \mathbf{v}_k is the sphere centered at \mathbf{c}_1 passing through \mathbf{v}_ℓ . Therefore, given any unit direction $d\mathbf{v}_k$, for a sufficiently small h , $\mathbf{v}_k + h d\mathbf{v}_k$ will no longer belong to the sphere, and the bisector of \mathbf{v}_k and \mathbf{v}_ℓ will not contain \mathbf{c}_1 . If the bisector contains another vertex \mathbf{c}_2 of $\Gamma_{\ell|F}$, we can apply the same principle. Therefore, for h small enough, the bisector of $\mathbf{v}_k + h d\mathbf{v}_k$ and \mathbf{v}_ℓ contains none of the vertices of $\Gamma_{\ell|F}$. Let \mathbf{x}_1 and \mathbf{x}_2 be the two intersections between the bisector of $\mathbf{v}_k + h d\mathbf{v}_k$ and \mathbf{v}_ℓ and the boundary of $\Gamma_{\ell|F}$, and let Ω'_k be the Voronoi cell of $\mathbf{v}_k + h d\mathbf{v}_k$. Since $[\mathbf{x}_1, \mathbf{x}_2]$ is the only edge depending on the position of \mathbf{v}_k ,

using Reynolds transport theorem, we obtain

$$\frac{\partial}{\partial \mathbf{v}_k} \left[\int_{\Gamma_{\ell|F} \cap \Omega'_k} \phi d\mathbf{x} \right] = \int_{\Gamma_{\ell|F} \cap \Omega'_k} \frac{\partial \phi}{\partial \mathbf{v}_k} d\mathbf{x} + \int_{[\mathbf{x}_1, \mathbf{x}_2]} \phi \frac{d\mathbf{x}}{d\mathbf{v}_k} \cdot \mathbf{n}_b d\mathbf{x},$$

where ϕ and its derivatives are evaluated at $(\mathbf{x}, \mathbf{V} + h d\mathbf{V})$, with $d\mathbf{V}$ being the perturbation of the samples which only moves \mathbf{v}_k in the direction $d\mathbf{v}_k$. Using [Theorem 1](#), when $\frac{\partial \phi}{\partial \mathbf{v}_k}$ is continuous, the first term in the sum above is continuous, and therefore we have

$$\lim_{h \rightarrow 0} \int_{\Gamma_{\ell|F} \cap \Omega'_k} \frac{\partial \phi}{\partial \mathbf{v}_k}(\mathbf{x}, \mathbf{V} + h d\mathbf{V}) d\mathbf{x} = \int_{\Gamma_{\ell|F} \cap \Omega_k} \frac{\partial \phi}{\partial \mathbf{v}_k}(\mathbf{x}, \mathbf{V}) d\mathbf{x}$$

For the second term, using [Equation 8](#) which does not depend on the parameterization of the edge, we can rewrite it

$$\|\mathbf{x}_1 - \mathbf{x}_2\| \int_0^1 \phi(\mathbf{x}(u), \mathbf{V} + h d\mathbf{V}) \frac{\mathbf{v}_k + h d\mathbf{v}_k - \mathbf{x}(u)}{\mathbf{n}_b \cdot (\mathbf{v}_k + h d\mathbf{v}_k - \mathbf{v}_\ell)} du \quad (19)$$

Since the inner part of the integral is bounded for all h in $(0, \rho)$ and for all $(\mathbf{x}_1, \mathbf{x}_2) \in F^2$, we can use Lebesgue dominated convergence theorem to move a limit into the integral and therefore obtain

$$\lim_{h \rightarrow 0} \int_{[\mathbf{x}_1, \mathbf{x}_2]} \phi(\mathbf{x}, \mathbf{V} + h d\mathbf{V}) \frac{d\mathbf{x}}{d\mathbf{v}_k} \cdot \mathbf{n}_b d\mathbf{x} = \int_{[\mathbf{c}_1, \mathbf{c}_2]} \phi(\mathbf{x}, \mathbf{V}) \frac{d\mathbf{x}}{d\mathbf{v}_k} \cdot \mathbf{n}_b d\mathbf{x}.$$

Combining these two limits, we obtain that whatever the displacement direction $d\mathbf{v}_k$

$$\begin{aligned} \lim_{h \rightarrow 0} \frac{\partial}{\partial \mathbf{v}_k} \left[\int_{\Gamma_{\ell|F} \cap \Omega'_k} \phi(\mathbf{x}, \mathbf{V} + h d\mathbf{V}) d\mathbf{x} \right] \\ = \int_{\Gamma_{\ell|F} \cap \Omega_k} \frac{\partial \phi}{\partial \mathbf{v}_k}(\mathbf{x}, \mathbf{V}) d\mathbf{x} + \int_{[\mathbf{x}_1, \mathbf{x}_2]} \phi(\mathbf{x}, \mathbf{V}) \frac{d\mathbf{x}}{d\mathbf{v}_k} \cdot \mathbf{n}_b d\mathbf{x}, \end{aligned}$$

which proves that the partial derivative of \mathcal{F} with respect to \mathbf{v}_k is continuous. Since all the partial derivatives of \mathcal{F} are continuous at \mathbf{V} , \mathcal{F} is continuously differentiable at \mathbf{V} . \square

Appendix B: Approximating piecewise smooth functions

In the form provided by Cortes *et al.* [\[CMB05\]](#), Reynolds transport theorem cannot be applied when f is not continuous. This property is also used in [Appendix A](#) to prove the continuity of our objective function \mathcal{F} . These results extend when f is only *piecewise* continuous. By definition, being piecewise continuous means that \mathcal{S} can be partitioned into a set of compact subset in which f is continuous and on the boundary of which f can be continuously extended. By segmenting \mathcal{S} between these compact sets, we obtain a set of meshes (with curved boundaries), and f is continuous on each of these meshes. We therefore fall back to handling meshes with boundaries with a continuous f . Since the discontinuities of f add edges to the mesh, discontinuities also appear in the objective function gradient when bisectors align with these new edges.