



On Asymptotically Optimal Meshes by Coordinate Transformation

Citation

Cañas, Guillermo D. and Steven J. Gortler. 2006. On asymptotically optimal meshes by coordinate transformation. In Proceedings of the 15th International Meshing Roundtable, September 17-20, 2006, Birmingham, Alabama, ed. International Meshing Roundtable and Philippe P. Pébay, 289-305. Berlin: Springer.

Published Version

http://dx.doi.org/10.1007/978-3-540-34958-7_17

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:2634286>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

On Asymptotically Optimal Meshes by Coordinate Transformation

Guillermo D. Cañas and Steven J. Gortler

Harvard University {gdiez|sjg}@cs.harvard.edu

Summary. We study the problem of constructing asymptotically optimal meshes with respect to the gradient error of a given input function. We provide simpler proofs of previously known results and show constructively that a closed-form solution exists for them. We show how the transformational method for obtaining meshes, as is, cannot produce asymptotically optimal meshes for general inputs. We also discuss possible variations of the problem definition that may allow for some forms of optimality to be proved.

1 Introduction

In this paper, we study “optimal triangular meshes for minimizing the gradient error”, as first described in the landmark paper by D’Azevedo and Simpson in [2].

In particular they consider the problem of approximating an input bivariate scalar function $f(x_1, x_2)$ by a piece-wise linear triangle mesh that interpolates f at its vertices. They are interested in approximations whose gradient errors are less than an input ϵ in an L^∞ sense, and are interested in the limit behavior of the meshing method as $\epsilon \rightarrow 0$.

D’Azevedo and Simpson [2] described the following simple approach. Given f ,

- Find a reparametrization of f (described by a mapping from the plane to the plane) that is “isotropic” in gradient error.
- Lay down a regular triangulation in the reparametrized domain.
- Map this triangulation back to the original (x_1, x_2) domain (straightening out any edges that have become curves).
- Sample f at the vertices of this triangulation to produce the piece-wise linear approximation.

D’Azevedo and Simpson then show the following

- They prove that a reparametrization that equidistributes gradient error exists, and can be calculated as the solution to a specific two-dimensional differential equation.
- They prove that if f is quadratic, then their method has at most 30% more triangles than any other triangulation that has L^∞ gradient error less than ϵ .
- They argue that the same 30% bound should hold *in the limit* for arbitrary f over any region where the determinant of H , the Hessian of f , is bounded away from 0.

This method is very appealing because it creates an approximation using a triangulation that has completely regular connectivity, a property that can be very useful for finite element applications. Moreover, the triangulation is the result of an explicit reparametrization step and so no optimization is required. Of course the reparametrization does require knowledge of f , so this may present somewhat of a bootstrapping problem, but in practice f may be approximated. Additionally, similar to some other techniques [2, 3, 6], this method only applies in regions where the determinant of H is bounded away from 0, and so for more complicated functions it may need to be used in conjunction with other global decomposition methods.

In this paper, we present the following new results

- We show that the isotropic reparametrization can be expressed in closed form simply as the gradient of f . No differential equation solving is necessary.
- We show that, in general, for a non-quadratic f the argument from [2] is not complete, and in fact, when there are no bounds on the anisotropy of H , then the amount of gradient error can be unbounded.
- We follow this with a discussion of some ideas we have explored to address the limitations of the basic algorithm.

Our closed form expression for the reparametrization greatly simplifies the method of [2]. It removes the need to solve a complicated differential equation (compare [2] equation 3.3). It also makes the existence proof of the reparametrization almost trivial (compare [2] section 7).

Unfortunately, we also show that for arbitrary f , the gradient error can be unbounded. Informally speaking, a right isosceles triangle in the equilibrating reparametrization will automatically have the “correct” aspect ratio as determined by the eigenvalues of H , and will be “correctly” stretched in the directions corresponding to the eigenvectors of H . But if the edges are not additionally aligned with the eigenvector directions, then the resulting triangulation can have large angles, arbitrarily close to 180 degrees.

For a quadratic function f , the eigenvector directions are spatially invariant, and so one can always rotate the domain to align its axes everywhere with the eigenvectors. For a non-quadratic function, this is generally not possible. These resulting triangles can thus exhibit arbitrary gradient error.

Negative results of this kind have appeared elsewhere. [5] shows that it is, in general, not possible to obtain meshes through the transformational method that satisfy the orientation and equidistribution property for any adaptation function that specifies the shape and orientation of elements at every point. Although ours is somewhat of a negative result, we hope that the readers will value the clarification of the strengths and weaknesses of the method described in [2].

2 Reparametrization

Consider the problem of approximating an input scalar function $f(x_1, x_2)$ defined over the plane with a piece-wise linear triangle mesh. The mesh exactly interpolates \mathbf{f} at the vertices and is linear inside triangles. Such a mesh defines a continuous piece-wise linear approximation of \mathbf{f} . We define the point-wise gradient error at each (x_1, x_2) to be the squared L^2 norm of the difference between the gradient of \mathbf{f} and the gradient of the mesh approximation. Because the mesh represents a piece-wise linear function, it's gradient is piece-wise constant, and can be evaluated only in the interior of triangles. This definition of gradient error follows [2].

Define the gradient error of the triangulation to be the maximum point-wise gradient error over the entire domain.

At every point in the plane, define the Hessian matrix of f with respect to the coordinates x_i , with entries

$$H_{ij} = \frac{\partial^2 f}{\partial x_i \partial x_j}$$

As done in [2], we restrict ourselves to a domain in the plane where H is non-degenerate (determinant bounded away from 0).

At each point, define the (positive definite) squared Hessian matrix $Q = H^T H$. As described in [2], this matrix is related to the gradient error for very small triangles. If, as in [2], we assume that for every small triangle, there is an interior (or very close) point x_0 where the triangle's gradient exactly matches the gradient of f , then the gradient error at any other point x interior to the triangle is approximately $H(x - x_0) = Hr$, where H is the Hessian matrix of the Taylor expansion of \mathbf{f} around x_0 , and $r = x - x_0$ is a very small displacement. The squared L^2 gradient error at a point is thus $r^T H^T H r = r^T Q r$ (plus higher order terms).

Under the transformational method it is our goal to find a reparametrization of the plane such that Q becomes the identity matrix.

In [2] it is shown that such an *isotropic* transformation always exists for any smooth \mathbf{f} , and they provide a differential equation that can be numerically integrated to find the solution. As discussed in [2], finding such a reparametrization has implications in finding triangulations with low gradient error.

2.1 A Closed Form Reparametrization

Here we present a simple closed form for a reparametrization of the plane such that Q becomes the identity matrix.

We interpret the matrix $Q = H^T H$ of section 2 as the coordinates of the $(0, 2)$ error tensor

$$q_{ab} = \sum_{ij} Q_{ij} (dx^i)_a (dx^j)_b$$

where the a and b subscripts are coordinate-free tensor place holders and the $(dx^i)_a$ make up the basis of covectors at each point [1].

Assume that we have reparametrized the plane using new coordinates

$$[\hat{x}_1(x_1, x_2), \hat{x}_2(x_1, x_2)]$$

We can then re-express the original tensor q as

$$q_{ab} = \sum_{ij} \hat{Q}_{ij} (d\hat{x}^i)_a (d\hat{x}^j)_b$$

with some new appropriate matrix \hat{Q} . Define the Jacobian matrix of the reparametrization at each point as

$$J_{ij} = \frac{\partial \hat{x}_i}{\partial x_j}$$

then, following the well known “basis change rule” [1], we have

$$\hat{Q} = J^{-T} Q J^{-1} \tag{1}$$

Note that this matrix \hat{Q} is *not* the square of a matrix with entries $\frac{\partial^2 f}{\partial \hat{x}_i \partial \hat{x}_j}$. Rather it represents the coordinates of the original tensor re-expressed in the new parametrization.

Define a parametrization (\hat{x}_1, \hat{x}_2) to be isotropic if $\hat{Q} = I$.

Define the gradient parametrization as

$$\hat{x}_i = \frac{\partial f}{\partial x_i} \tag{2}$$

Theorem 1. *Given a bivariate function $f(x_i)$ in a region where its Hessian matrix is non-singular, then the gradient parametrization is an isotropic parametrization.*

Proof. The entries of the Jacobian of this parametrization are

$$J_{ij} = \frac{\partial}{\partial x_j} \left(\frac{\partial f}{\partial x_i} \right) = \frac{\partial^2 f}{\partial x_j \partial x_i} = H_{ij}$$

Under this parametrization, the matrix \hat{Q} can be computed as

$$\hat{Q} = J^{-T} Q J^{-1} = J^{-T} (H^T H) J^{-1} = H^{-T} (H^T H) H^{-1} = I$$

■

This proves that an isotropic parametrization exists, and that it can be computed in closed-form using equation (2).

Theorem 2. The isotropic parametrization of equation (2) is unique up to rigid isometry.

Proof. Consider any isotropic parametrization $\tilde{x}(x)$ that is different from the gradient parametrization $\hat{x}(x)$. Because the Jacobian of $\hat{x}(x)$ is the Hessian (H) of f , and H is everywhere non-singular, then $\hat{x}(x)$ is invertible. We can always write the new isotropic parametrization as $\tilde{x} = s \circ \hat{x}$, where $s = \tilde{x} \circ \hat{x}^{-1}$. In a certain sense, s tells us how \tilde{x} differs from \hat{x} .

If K is the Jacobian of $\tilde{x}(x)$, then $S \equiv KH^{-1}$ is the Jacobian of s . From equation (1) and the fact that \tilde{x} is an isotropic parametrization, we know that it must satisfy

$$\hat{Q} = K^{-T} Q K^{-1} = K^{-T} (H^T H) K^{-1} = I$$

or in other words

$$(KH^{-1})^T (KH^{-1}) = S^T S = I$$

Which shows that S is everywhere orthogonal. It follows that, since s is a mapping between two planar domains with an everywhere orthogonal Jacobian, s must be a rigid isometry [7]. ■

It is also a tedious but straightforward calculation to show that the gradient parametrization, in fact, satisfies the differential equation given in [2] (equations 3.3 to 3.6).

We first note that, in the above proofs, there is no explicit mention of the dimensionality of the input domain, suggesting that the above is a proof that the gradient map is the solution to any equivalent problem posed in an arbitrary-dimensional Euclidean space.

Corollary 1. *Given an input function $f : \mathbb{R}^n \rightarrow \mathbb{R}$, in a region where its Hessian is non-singular, the gradient parametrization of equation (2) is an isotropic parametrization. This isotropic parametrization is unique up to rigid isometry.*

We can conclude that, given an input with a non-singular Hessian, there always exists a mapping that can transform the original domain into one in which the gradient error is isotropically distributed, and that this mapping is always the gradient map composed with some rigid isometry.

2.2 Alignment

For purposes of error analysis, it will later be important to study the relation between the orientations of edges of a triangulation in the original and reparametrized spaces. We present this analysis here to establish some notation and concepts that will be useful through later sections.

We can study the relation between the orientations of edges in the original and transformed domains for meshes that are increasingly fine by looking at simple differential properties of a reparametrization $m : \mathbb{R}^n \rightarrow \mathbb{R}^n$. Given a triangle edge starting at \mathbf{x} and oriented in direction \mathbf{v} , it will be oriented in direction $\mathbf{w} = m(\mathbf{x} + \mathbf{v}) - m(\mathbf{x})$ when transformed by m . We can see that, in the limit as triangles become smaller and smaller, and their edges are equally shrunk, an edge with orientation \mathbf{v} will map onto an edge oriented in the direction of $\mathbf{w} = D_{\mathbf{v}}m$ (the derivative of m in the direction \mathbf{v}) in the reparametrized domain.

Given an input function f with an everywhere non-singular Hessian, **Define** an *aligned isotropic* parametrization to be an isotropic parametrization that maps the eigenvector directions of H onto the coordinate directions of the isotropic domain.

From our knowledge of the properties of isotropic parametrizations, we can easily compute the form that an aligned isotropic parametrization has. If the (generally spatially varying) Hessian of the input is $H = R^T \Lambda R$, where R is orthogonal, and Λ is diagonal, then, as proved in section 2.1, any isotropic parametrization m will have a Jacobian $K = SH = SR^T \Lambda R$, where S is a constant orthogonal matrix.

This parametrization will transform the i -th eigenvector \mathbf{v}_i of H into

$$\mathbf{w}_i = D_{\mathbf{v}_i}m = K\mathbf{v}_i = SR^T \Lambda R\mathbf{v}_i = (SR^T) \cdot (\lambda_i \hat{e}_i) \quad (3)$$

where λ_i is the eigenvalue corresponding to \mathbf{v}_i and \hat{e}_i is the i -th coordinate direction. By definition, an aligned isotropic parametrization is one for which, for all i , this \mathbf{w}_i is in the direction of \hat{e}_i . From equation (3) and the fact that SR^T is an orthogonal matrix, we conclude that for m to be aligned it must be that $SR^T = I$, that is $S = R$. In conclusion, we see that an aligned isotropic parametrization is an isotropic parametrization that is the composition of the gradient parametrization with an isometry with Jacobian S , where at every point it is $S = R$.

Because, as proved in section 2.1, the orthogonal matrix S is constant, we can state that an aligned isotropic parametrization only exists for input functions f for which R is constant (e.g. a quadratic f). In conclusion, for general inputs f for which R is spatially varying, an isotropic reparametrization exists, but an aligned isotropic one doesn't. This fact will prove to be crucial in later sections.

3 Asymptotically Optimal Meshing

We are now in a good position to analyze the problem of asymptotically optimal meshing with respect to gradient error as defined above. Consider a mesh such that every triangle in the mesh has gradient error (as defined in section 1) bounded by a constant ϵ . For any given mesh and bound ϵ , we define the *efficiency* of the mesh to be the average area of its triangles. An optimal mesh is one that has the highest efficiency among all others for a given ϵ bound, because it must have the minimum number of triangles when meshing any finite domain.

Given a mesh \mathcal{T} with an error bound of ϵ and efficiency $\nu(\mathcal{T})$, its *competitiveness ratio* is defined as $\nu(\mathcal{T})/\nu_o(\epsilon)$, where $\nu_o(\epsilon)$ is the optimal efficiency for this error bound. Note that for a domain which is a finite area of the plane (x_1, x_2) , we can also measure the competitiveness ratio using the “average area” in a reparametrized domain (\hat{x}_1, \hat{x}_2) . This is true because the total area before and after the reparametrization is simply related by a constant.

A meshing algorithm \mathcal{A} that can take any ϵ and produce a mesh $\mathcal{T} = \mathcal{A}(\epsilon)$ with an error bound of ϵ is said to be asymptotically optimal to within a constant λ if $\lim_{\epsilon \rightarrow 0} \frac{\nu(\mathcal{A}(\epsilon))}{\nu_o(\epsilon)} \geq \lambda$.

3.1 Transformational Meshes

We describe here the process of obtaining a mesh from a transformation, and the subsequent steps that we take to analyze the error of its triangles. We first assume that the transformation, or reparametrization, that we apply to the original domain where the input \mathbf{f} is defined is an isotropic transformation, and therefore is the gradient map composed with some isometry. As before, we call the original domain x and the isotropic one \hat{x} .

For any given value of the error threshold ϵ , we seek to find a triangulation with squared gradient error bounded by ϵ . We first lay out a regular grid of right-isosceles triangles in the \hat{x} -domain, as in figure 1 (left). We then map the vertices of these triangles to the x -domain and connect them with straight edges using their original connectivity, as shown in figure 1 (middle). This defines the transformational mesh. The function f is then sampled at the vertices of the transformational mesh, and approximated as a piece-wise linear function.

To analyze the error of the transformational mesh’s triangles we take another further step, and transform the mesh back to the isotropic domain. Consider the highlighted triangle in the x -domain of figure 1 (middle), with vertices x_1 , x_2 , and x_3 . When we map the triangle back onto the \hat{x} -domain we obtain the highlighted triangle $\triangle \hat{x}_1 \hat{x}_2 \hat{x}_3$ of figure 1 (right). Notice that, although the vertices of this triangles lie on a uniform grid, its edges are in general curved, since the gradient map is in general a non-linear transformation. We can compute the gradient of the triangle $\triangle x_1 x_2 x_3$, which is simply




Fig. 1. A regular grid of right-isosceles triangles in the reparametrized domain (left). The corresponding set of triangles shown in the original domain (a transformational mesh) (middle). The marked triangle transformed back to the reparametrized domain (right).

the gradient of the piece-wise linear approximation formed by $\Delta x_1 x_2 x_3$. This gradient can be identified as some point $\hat{\mathbf{x}}^*$ in the \hat{x} -domain, as shown in figure 1 (right). This is so because the \hat{x} -domain really represents (up to isometry) the gradient space of the input function, and so the gradient $\hat{\mathbf{x}}^*$ will correspond to some location in the \hat{x} -domain.

Once this setup is complete, the task of measuring gradient error is simple. In the x -domain, the squared gradient error between any point in the interior of a triangle and the piece-wise linear approximating mesh is the squared distance between the gradient of \mathbf{f} at that point and the gradient of the triangle that the point lies in. In other words, we can measure this error in the \hat{x} -domain by simply taking squared euclidean distances between any point interior to a transformed triangle and its gradient $\hat{\mathbf{x}}^*$. Finally, we point out that composing the gradient parametrization with an isometry does not change this reasoning since, by definition, isometries do not change distances between transformed points. Thus the above analysis applies to any isotropic parametrization.

3.2 Quadratic Input

We first analyze the behavior of the above construction when applied to a simple quadratic input \mathbf{f} . The relevance of this simpler case is that, as pointed out in [2], the behavior of an arbitrary input is very similar to that of the quadratic case when we look at meshes in the limit as they become finer and finer.

Consider an input function \mathbf{f} with constant Hessian. This Hessian H is symmetric real and can be decomposed as $H = R^T \Lambda R$, where Λ is diagonal and R is a rotation by θ radians. Let us refer to the eigenvalues of H as λ_1 and λ_2 . Consider for analysis a right-isosceles triangle in the \hat{x} domain with coordinates $\hat{x}_1 = \hat{o} + \sqrt{\frac{\epsilon}{2}}(-1, -1)$, $\hat{x}_2 = \hat{o} + \sqrt{\frac{\epsilon}{2}}(1, -1)$, and $\hat{x}_3 = \hat{o} + \sqrt{\frac{\epsilon}{2}}(-1, 1)$

(because it is expressed in generic form as a function of the coordinates of its circumcenter \hat{o} , this analysis applies to any triangle in a regular grid of right-isosceles triangles.) We use an isotropic parametrization \hat{x} in general form, that is, computed as the gradient parametrization composed by an isometry with Jacobian S , where S is a rotation by α radians. Note that, because the gradient map of a quadratic function is a linear map, both this triangle and its corresponding triangle in the x -domain have straight edges. In order to analyze the error inside this triangle, we compute its gradient ∇f^* (the gradient of its supporting plane) and the location of this gradient in the \hat{x} -domain: $\hat{x}^* = S\nabla f^*$. We can show (see appendix) that \hat{x}^* is such that

$$\|\hat{x}^* - \hat{o}\| = \sqrt{\epsilon} |\sin[2(\theta - \alpha)]\mu| \tag{4}$$

with

$$\mu = \frac{1}{2\sqrt{2}}(\lambda_1 - \lambda_2) \left[\frac{\lambda_1}{\lambda_2}(1 + \sin(2\theta)) + \frac{\lambda_2}{\lambda_1}(1 - \sin(2\theta)) \right] \tag{5}$$

From equation (4) we can conclude that if $\theta = \alpha$, that is, if $S = R$ and therefore if the isotropic reparametrization is aligned, as defined in section 2.2, then $\hat{x}^* = \hat{o}$. In this case, because the error of a triangle is the maximum distance between any of its interior points and \hat{x}^* , we can conclude that the error of this triangle is exactly ϵ . (Recall that since f is quadratic, the isotropic image of $\Delta x_1 x_2 x_3$ is a triangle with straight edges.)

Clearly, as was argued in [2], in any optimal triangulation, the isotropic image of each triangle cannot have an area greater than that of an equilateral triangle circumscribed by a circle of radius $\sqrt{\epsilon}$ (otherwise, it could not have gradient error bounded by ϵ). Each of our right-isosceles triangles is only 23% smaller than such an equilateral triangle, so the average area (measured in the isotropic domain) of the triangles obtained by this algorithm can be no smaller than 77% of the average area of the triangles of an optimal triangulation. Therefore if the reparametrization is aligned isotropic, we conclude, in agreement with [2], that the induced transformational mesh is asymptotically optimal.

We can also compute a lower bound on the approximating error (E) for any right-isosceles triangle. In particular, for any point \hat{x} interior to the triangle we have

$$\begin{aligned} E(\hat{x}) &= \|\hat{x} - \hat{x}^*\|^2 = \|\hat{x} - \hat{o} - (\hat{x}^* - \hat{o})\|^2 \\ &\geq \max\{0, \|\hat{x}^* - \hat{o}\| - \|\hat{x} - \hat{o}\|\}^2 \\ &\geq \max\{0, \|\hat{x}^* - \hat{o}\| - \sup\{\|\hat{x} - \hat{o}\|\}\}^2 \\ &\geq \max\{0, \sqrt{\epsilon} |\sin[2(\theta - \alpha)]\mu| - \sqrt{\epsilon}\}^2 \\ &= \epsilon \cdot \max\{0, |\sin[2(\theta - \alpha)]\mu| - 1\}^2 \end{aligned} \tag{6}$$

We can notice from equation (6) that if $\theta \neq \alpha$, that is, if $S \neq R$ (the reparametrization is not aligned), then the gradient error in a triangle can

be arbitrarily large if we are not to place any restrictions on the amount of anisotropy of f .

Intuitively, the amount of gradient *variation* of f inside such a triangle is bounded by ϵ , but in the non-aligned case, the gradient of the piece-wise linear approximation can still be significantly different from the gradient of f (everywhere) in the triangle.

In figure 2 we show two triangulations resulting from isotropic parametrizations. On the left, we show the point-wise gradient error obtained when using the aligned parametrization. In this case, the triangles all have small (90° or less) angles, they have zero error at their circumcenter (at the midpoint of one of their edges), and error bounded by ϵ elsewhere. On the right, we show the result using a non-aligned parametrization. In this case, the triangles can have arbitrarily large error, they might not have zero error anywhere (they might actually not interpolate the gradient), and they often have large angles. This link between large gradient error and large angles has been pointed out before (see [4] for example).

In [2] an aligned isotropic reparametrization for the quadratic case is computed ([2] equation 2.8), and in [2] theorem 2.1, it is proved that this will result in an optimal triangulation. As we have shown, there are in fact many isotropic reparametrizations, but only the one that is aligned has guarantees that it leads, for any quadratic input, to an asymptotically optimal mesh. This distinction is not made explicit in [2]. This leads them to incorrectly conclude in their section 3 (the non-quadratic case), that the isotropic parametrization they find (which is generally not everywhere aligned), will result in an asymptotically optimal triangulation.




Fig. 2. Gradient approximation error graphed for a transformational mesh derived from an aligned isotropic reparametrization (left), and a non-aligned isotropic one (right). The bar on the far right shows the error scale (white is zero error). The input function is $f(x_1, x_2) = 0.156(x_1)^2 + 0.318(x_2)^2 - 0.281x_1x_2$.

3.3 Arbitrary Input

For non quadratic input, the eigenvector directions of $H = R^T \Lambda R$, and therefore R , will in general be spatially varying. As a result, for an isotropic reparametrization with Jacobian SH where S is orthogonal, (except for a curve of measure 0) we will generally have $S \neq R$, which corresponds to $\theta \neq \alpha$ in equation 6. In these regions, if $\lambda_1 \gg \lambda_2$ (f is highly anisotropic), then for any regular right-isosceles triangulation of the isotropic reparametrized domain, the gradient error will be accordingly large, and the triangles will have to be made much smaller than optimal. If no bounds are placed on the anisotropy of f , then to meet a gradient error threshold the triangles may be arbitrarily small, and an optimality bound cannot in general be met.




Fig. 3. Point-wise error of a transformational triangulation (left), and worst per-triangle error of the same triangulation (right). The bar on the far right shows the error scale (white is zero error). The lowest worst error is in the triangles in a vertical band around the center of each triangulation, where the eigenvectors of H are better aligned with the coordinate directions in the isotropic reparametrization. The input function is $f(x_1, x_2) = 1/\sqrt{(x_1)^2 + (x_2)^2}$. Shown is a region slightly away from the singularity.

One can of course still obtain upper bounds on the error, if one imposes an anisotropy bound on f , but this is a much weaker result than hoped for, and much less than what can be achieved in the quadratic case, where asymptotic optimality is guaranteed for an arbitrary quadratic input.

In practice we have seen that the error can grow quite large even when the amount of anisotropy is not very extreme. In figure 3 we show a transformational mesh obtained from an isotropic reparametrization. The lowest errors are found in a vertical band around the center of each image, where the alignment happens to be best (S is closest to R).

3.4 Relaxing Continuity




Fig. 4. A non-conforming triangle mesh that is C^0 continuous only at the edge midpoints in the isotropic domain (left). The construction used in the appendix for showing C^0 continuity at midpoints and bounded error (right).

We can attempt to relax the problem definition in a way that might make it easier to obtain asymptotically optimal meshes. In particular, we can consider a variation of the problem where we do not sample f at vertices, but directly interpolate both f and its gradient at some interior point of a subset of the mesh’s triangles.

Clearly if we relax the problem definition and allow the mesh to not have C^0 continuity, then it is possible to obtain asymptotically optimal meshes in the limit. We can lay out a uniform grid of triangles in isotropic space, and simply treat each triangle separately. We then require that each triangle exactly interpolate the function f and its gradient at its circumcenter. The error in each triangle then will be bounded by ϵ , the resulting mesh is optimally efficient in the limit. Since this method is simply interpolating f and its gradient at its center, we can just as easily use a uniform grid of equilateral triangles.

In practice a mesh that doesn’t even have C^0 continuity might not be very useful. We can do slightly better by constructing the “non-conforming” mesh of figure 4, which will be guaranteed to be C^0 at the edges’ midpoints. In this mesh of equilateral triangles (in isotropic space), the input f and its gradient are sampled at the circumcenter of every marked triangle. This determines the piece-wise linear approximation \bar{f} at triangles that are marked. We construct \bar{f} at all other triangles by reading the values of \bar{f} at the midpoints of edges of marked triangles. We can see that every edge in the triangulation is incident to exactly one marked triangle, and so this construction exactly determines \bar{f} without over-constraining it.

In appendix B, we show that in the quadratic case, this produces asymptotically optimal meshes when any isotropic reparametrization is used, even an unaligned one. Because in the non-quadratic case we can always compute

an isotropic parametrization and alignment is not necessary, and because in the limit the approximation's behavior is dominated by its low order behavior, we conjecture that a formal proof of its optimality can be constructed.

Again, the obvious drawback is that this is a non-conforming triangle mesh that is only guaranteed to be C^0 at edge midpoints.

4 Asymptotically Optimal Triangulations

Although we have shown that a transformational mesh, as defined, cannot produce guarantees on asymptotical optimality for an arbitrary input, our analysis does provide a deeper understanding of the problem of asymptotically optimal meshing.

We have also shown how a regular grid of right-isosceles triangles in an isotropic reparametrization has error bounded by ϵ if all triangles are perfectly aligned (that is, if their edges in the original parametrization are aligned with the eigenvectors of H), and that their efficiency is ϵ : only 23% lower than the maximum attainable efficiency. Therefore such a triangulation is guaranteed to be asymptotically optimal to within at least 77%.

In practice, obtaining such a triangulation may prove too hard, and we can obviously allow some flexibility by permitting the triangulation to include non-aligned triangles that have an error bounded by ϵ , and an area possibly much smaller than the maximum attainable efficiency, so long as the proportion of these triangles to the total goes to zero when we make the triangulation increasingly fine. This approach may provide a future avenue for obtaining provably asymptotically optimal meshes. We plan to explore this possibility in future work.

References

1. Robert M. Wald (1984) *General Relativity*. University of Chicago Press, Chicago.
2. E. F. D'Azevedo and R. B. Simpson (1991) On Optimal Triangular Meshes for Minimizing the Gradient Error. *Numerische Mathematik*. 59, 321–348
3. E. F. D'Azevedo (1991) Optimal Triangular Mesh Generation by Coordinate Transformation. *SIAM Journal on Scientific and Statistical Computing*. 12, 4, 755–786.
4. J. R. Shewchuk (2002) What Is a Good Linear Finite Element? - Interpolation, Conditioning, Anisotropy, and Quality Measures. Eleventh International Meshing Roundtable. 115–126.
5. W. Huang (2005), Anisotropic mesh adaptation and movement. Lecture notes for Peking Univ. Workshop on Adaptive mesh methods.
6. L. Chen and J. Xu (2004) Optimal Delaunay Triangulations. *Journal of Computational Mathematics*. 22, 2, 299–308.
7. M. Do Carmo (1976) *Differential Geometry of Curves and Surfaces*. Prentice Hall, New Jersey.

Appendix A

Suppose that we are given an input quadratic function \mathbf{f} with Hessian $H = R^T \Lambda R$, where, as in section 3.2, R is a rotation by θ radians, and Λ has eigenvalues λ_1 and λ_2 . We consider an isotropic parametrization \hat{x} of the x -domain that is the composition of the gradient parametrization with an isometry having orthogonal Jacobian S , where S is a rotation of the plane by α radians. We first compute the expression of the gradient of a triangle that, in the \hat{x} -domain, has vertices with coordinates $\hat{x}_1 = \sqrt{\frac{\epsilon}{2}}(-1, -1)$, $\hat{x}_2 = \sqrt{\frac{\epsilon}{2}}(1, -1)$, and $\hat{x}_3 = \sqrt{\frac{\epsilon}{2}}(-1, 1)$. This triangle can be mapped to the x -domain and considered as a piece-wise linear element. As such, it will have a gradient, which we then transform to the \hat{x} -domain by applying the rotation S . The final coordinates of this transformed gradient \hat{x}^* in the \hat{x} -domain can be used to derive error bounds for the triangle. We then show that if we were to translate all the points of this triangle in the \hat{x} -domain by a vector \hat{o} we would get the same result, that is, the value of \hat{x}^* is the same as before except that it is also translated by \hat{o} .

If the input function is an arbitrary quadratic function $f(x) = c + g^T x + \frac{1}{2}x^T Hx$, where c is an arbitrary constant and g an arbitrary vector, then we can write the relation between the x and the \hat{x} domains. Because the parametrization is by definition $\hat{x} = S\nabla f(x)$, and $\nabla f(x) = g + Hx$, then we can say that

$$x = H^{-1}(S^T \hat{x} - g) \quad (7)$$

Because the triangle $\triangle x_1 x_2 x_3$ linearly approximates f in the x -domain, we can write

$$f^* - f_1 = u \cdot (f_2 - f_1) + v \cdot (f_3 - f_1)$$

where $f^*(x)$ is the linear approximation, $f_i = f(x_i)$, and u and v are (the first two) local barycentric coordinates of the triangle. The scalars (u, v) are such that at any point x we can write

$$x - x_1 = P \begin{bmatrix} u \\ v \end{bmatrix}$$

$$P = [x_2 - x_1 \quad x_3 - x_1]$$

We can write the gradient of f^* as

$$\frac{\partial f^*}{\partial x_i} = \frac{\partial f}{\partial u} \frac{\partial u}{\partial x_i} + \frac{\partial f}{\partial v} \frac{\partial v}{\partial x_i}$$

which can be rewritten as

$$\nabla f^* = P^{-1} \begin{bmatrix} f_2 - f_1 \\ f_3 - f_1 \end{bmatrix} \quad (8)$$

We now compute the two terms of the right-hand side of equation (8).

Because $x_2 - x_1 = H^{-1}S^T(\hat{x}_2 - \hat{x}_1)$ and $x_3 - x_1 = H^{-1}S^T(\hat{x}_3 - \hat{x}_1)$, and from the definition of P and of x_1, x_2 , and x_3 we find that $P = \sqrt{2\epsilon}H^{-1}S^T$ and

$$P^{-1} = (\sqrt{2\epsilon})^{-1}SH$$

On the other hand, the expressions for $f_2 - f_1$ and $f_3 - f_1$ simplify to

$$\begin{aligned} f_2 - f_1 &= \frac{\epsilon}{4} \left(\begin{bmatrix} 1 \\ -1 \end{bmatrix}^T SH^{-1}S^T \begin{bmatrix} 1 \\ -1 \end{bmatrix} - \begin{bmatrix} -1 \\ -1 \end{bmatrix}^T SH^{-1}S^T \begin{bmatrix} -1 \\ -1 \end{bmatrix} \right) \\ f_3 - f_1 &= \frac{\epsilon}{4} \left(\begin{bmatrix} -1 \\ 1 \end{bmatrix}^T SH^{-1}S^T \begin{bmatrix} -1 \\ 1 \end{bmatrix} - \begin{bmatrix} -1 \\ -1 \end{bmatrix}^T SH^{-1}S^T \begin{bmatrix} -1 \\ -1 \end{bmatrix} \right) \end{aligned}$$

If we call $B = SHS^T$, then we know that $\det(B) = \det(H)$, and that B can be decomposed as $B = (RS^T)^T \Lambda (RS^T)$ where RS^T is a rotation of the plane by $\theta - \alpha$ radians. In particular, the off-diagonal entry of B is $B_{12} = \frac{1}{2} \sin(\theta - \alpha)(\lambda_1 - \lambda_2)$, and we can also verify that $B_{12}^{-1} = -B_{12}/\det(B) = -B_{12}/\det(H)$.

This allows us to simplify the above equations into

$$\begin{aligned} f_2 - f_1 &= \epsilon B_{12}/\det(H) \\ f_3 - f_1 &= \epsilon B_{12}/\det(H) \end{aligned}$$

From this we can compute

$$\begin{aligned} \nabla f^* &= P^{-1} \begin{bmatrix} f_2 - f_1 \\ f_3 - f_1 \end{bmatrix} \\ &= \sqrt{\epsilon} B_{12} \frac{SH}{\det(H)} \begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix} \end{aligned}$$

Because $\hat{x}^* = S\nabla f^*$, then

$$\hat{x}^* = \frac{\sqrt{\epsilon}}{2} \sin[2(\theta - \alpha)(\lambda_1 - \lambda_2)] \frac{S^2 H}{\det(H)} \begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix} \quad (9)$$

We can compute the norm of \hat{x}^* by taking into account that S^2 is an orthogonal matrix and therefore

$$\|\hat{x}^*\| = \frac{\sqrt{\epsilon}}{2\sqrt{2}} |\sin[2(\theta - \alpha)](\lambda_1 - \lambda_2)| \left[\frac{\lambda_1}{\lambda_2} (1 + \sin(2\theta)) + \frac{\lambda_2}{\lambda_1} (1 - \sin(2\theta)) \right] \quad (10)$$

We now consider the more general case of an arbitrary triangle $\mathcal{T} = \triangle \hat{x}'_1 \hat{x}'_2 \hat{x}'_3$ in the \hat{x} -domain that is the translation of the above $\triangle \hat{x}_1 \hat{x}_2 \hat{x}_3$ by an arbitrary vector \hat{o} in the \hat{x} -domain. Where in this notation it is $\hat{x}' = \hat{x} + \hat{o}$. In this case we want to compute $\hat{x}'^* - \hat{o}$ and its norm, where \hat{x}'^* is the gradient of triangle $\triangle x'_1 x'_2 x'_3$ transformed by S . We can write the relation between points in the \hat{x} and x -domains as

$$x = H^{-1}(S^T \hat{x}' - g) = H^{-1}(S^T \hat{x} + S^T \hat{o} - g) = H^{-1}[S^T \hat{x} - (g - S^T \hat{o})] \quad (11)$$

We now consider the problem of approximating not f but the function $\bar{f}(x) = f(x) - \hat{o}^T Sx$. Because the gradient of \bar{f} at the origin is $g - S^T \hat{o}$, we can see that equation (11) is formally the same as (7) if we are approximating \bar{f} as opposed to f . We can apply the same analysis as above and find that, because \bar{f} and f have the same Hessian, the approximated gradient of \bar{f} is

$$\nabla \bar{f}^* = \frac{\sqrt{\epsilon}}{2} \sin[2(\theta - \alpha)(\lambda_1 - \lambda_2)] \frac{SH}{\det(H)} \begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}$$

But because a linear approximation that is based on interpolating a function at three points reproduces linear functions, and $f = \bar{f} + \hat{o}^T Sx$, we find that if we had been approximating f we would've obtained that $\nabla f^* = \nabla \bar{f}^* + S^T \hat{o}$. We are now interested in computing $\hat{x}'^* - \hat{o}$, which we can obtain by multiplying the expression for ∇f^* by S , and subtracting \hat{o}

$$\hat{x}'^* - \hat{o} = \frac{\sqrt{\epsilon}}{2} \sin[2(\theta - \alpha)(\lambda_1 - \lambda_2)] \frac{S^2 H}{\det(H)} \begin{bmatrix} 1/\sqrt{2} \\ 1/\sqrt{2} \end{bmatrix}$$

whose norm is

$$\|\hat{x}'^* - \hat{o}\| = \frac{\sqrt{\epsilon}}{2\sqrt{2}} |\sin[2(\theta - \alpha)](\lambda_1 - \lambda_2) \left[\frac{\lambda_1}{\lambda_2} (1 + \sin(2\theta)) + \frac{\lambda_2}{\lambda_1} (1 - \sin(2\theta)) \right]|$$

Which has the same form as equations (9) and (10).

Appendix B

We show here that the triangulation shown in figure 4 (left) has gradient error bounded by ϵ and has \mathcal{C}^0 continuity at edge midpoints. Given an input function that is an arbitrary quadratic function $f(x) = c + g^T x + \frac{1}{2} x^T Hx$, where c is an arbitrary constant and g an arbitrary vector, then we can write the relation between the x and the \hat{x} domains. Because the parametrization is, by definition, $\hat{x} = S\nabla f(x)$, and $\nabla f(x) = g + Hx$, we can say that

$$x = H^{-1}(S^T \hat{x} - g) \quad (12)$$

Consider the triangles shown in figure 4 (right) in the isotropic domain \hat{x} . Their circumcenters are \hat{o}_i and the midpoint of their shared edges are $\hat{m}_i = \frac{1}{2}\hat{o}_1 + \frac{1}{2}\hat{o}_i$, $i \in \{2, 3, 4\}$ (their counterparts in the original domain are o_i and m_i). For those triangles that are "marked", the input f and its gradient

are sampled at their circumcenter \hat{o}_i , $i \in \{2, 3, 4\}$, and thus their gradient error is ϵ . This sampling completely determines the piece-wise linear approximation f^* of f inside marked triangles. At the midpoints \hat{m}_i of edges, we sample f^* from the single marked triangle incident to the edge. Setting the values of f^* at \hat{m}_2 , \hat{m}_3 , and \hat{m}_4 then completely determines the piece-wise linear approximation at \mathcal{T}_1 . \mathcal{C}^0 continuity at the midpoints \hat{m}_i is guaranteed by construction. To prove that the gradient error at \mathcal{T}_1 is bounded by ϵ we simply show that the construction above will produce the same values of f^* at the midpoints \hat{m}_i (and therefore the same linear approximation inside \mathcal{T}_1) as if we had sampled f and ∇f at the circumcenter of \mathcal{T}_1 .

For triangle \mathcal{T}_i , $i \in \{2, 3, 4\}$, we can compute the sampled value $f^*(m_i) = f(o_i) + \nabla f(o_i)^T(m_i - o_i)$. Similarly, for \mathcal{T}_1 , we can compute the value that f^* would have at \hat{m}_i if we had obtained it by sampling f and ∇f at o_1 : $\bar{f}^*(m_i) = f(o_1) + \nabla f(o_1)^T(m_i - o_1)$. We now only have to prove that $f^*(m_i) = \bar{f}^*(m_i)$ for $i \in \{2, 3, 4\}$.

We first compute $f(o_i)$, $i \in \{1, 2, 3, 4\}$. From (12) we know that $o_i = H^{-1}(S^T \hat{o}_i - g)$, and so

$$\begin{aligned} f(o_i) &= c + g^T H^{-1}(S^T \hat{o}_i - g) + \frac{1}{2}(\hat{o}_i^T S - g^T)H^{-1}(S^T \hat{o}_i - g) \\ &= c - g^T H^{-1}g + \frac{1}{2}g^T H^{-1}g + \frac{1}{2}\hat{o}_i^T S H^{-1}S^T \hat{o}_i \end{aligned} \quad (13)$$

Because (from our definition of the reparametrization) $\nabla f(o_i) = S^T \hat{o}_i$, we can write

$$\begin{aligned} f^*(m_i) &= f(o_i) + \nabla f(o_i)^T(m_i - o_i) \\ &= c - g^T H^{-1}g + \frac{1}{2}g^T H^{-1}g + \frac{1}{2}\hat{o}_i^T S H^{-1}S^T \hat{o}_i + \hat{o}_i^T S H^{-1}S^T(\hat{m}_i - \hat{o}_i) \\ &= c - g^T H^{-1}g + \frac{1}{2}g^T H^{-1}g + \frac{1}{2}\hat{o}_i^T S H^{-1}S^T \hat{o}_i + \hat{o}_i^T S H^{-1}S^T\left(\frac{1}{2}\hat{o}_i + \frac{1}{2}\hat{o}_1 - \hat{o}_i\right) \\ &= c - g^T H^{-1}g + \frac{1}{2}g^T H^{-1}g + \frac{1}{2}\hat{o}_i^T S H^{-1}S^T \hat{o}_1 \end{aligned} \quad (14)$$

While on the other hand

$$\begin{aligned} \bar{f}^*(m_i) &= f(o_1) + \nabla f(o_1)^T(m_i - o_1) \\ &= c - g^T H^{-1}g + \frac{1}{2}g^T H^{-1}g + \frac{1}{2}\hat{o}_1^T S H^{-1}S^T \hat{o}_1 + \hat{o}_1^T S H^{-1}S^T(\hat{m}_i - \hat{o}_1) \\ &= c - g^T H^{-1}g + \frac{1}{2}g^T H^{-1}g + \frac{1}{2}\hat{o}_1^T S H^{-1}S^T \hat{o}_1 \end{aligned} \quad (15)$$