Erica Schwarz
3/10/16
CIS II: Seminar Paper Critical Review
Group 12

Critical Review: Mesh Simplification and Measures of Error

Introduction

Project Review

Our project is "Validating and Improving Single-Stage Cranioplasty Prosthetics with Ground Truth Models". The aim of this project is to use patient-specific ground truth models of cranioplasty defects to determine accuracy, evaluate robustness, and improve upon patient registration of a process that segments the defect wall from a 3D scanned point cloud of the defect. If successful, this project will pave the way for accurate single-stage cranioplasty operations that will reduce surgical time and the risk of infection due to an ill-fitting cranioplasty implant.

Paper Selection

The following are the two papers I reviewed:

Hoppe, Hugues, et al. "Mesh optimization." Proceedings of the 20th annual conference on Computer graphics and interactive techniques. ACM, 1993.

Aspert, Nicolas, Diego Santa Cruz, and Touradj Ebrahimi. "MESH: measuring errors between surfaces using the Hausdorff distance." ICME (1). 2002.

The first paper, Hoppe et al., discusses a mesh optimization method that has applications in mesh simplification and segmentation. The second paper is shorter and discusses a robust method for gauging surface-to-surface error.

Significance

The Hoppe et al. paper could provide a method for addressing some of the problems we face with the current method of finding the defect: handling the dense 3D point cloud obtained from the scan of the defect. There are thousands of points in a single defect scan, which slows computation time, and the mesh representation of these points tends to be noisy, with surface normal pointing in very disparate directions even within a localized area. It would be ideal to work with a simplified representation of this mesh. Given the point cloud density, it is reasonable to claim that we can significantly reduce the number points without having a large effect on the overall accuracy of the mesh. This would allow for faster manipulation of the mesh and also reduce the noisy variation of surface normals in dense areas (this is important for the method used to identify the parts of the mesh that are the defect wall). In addition, the segmentation applications of this approach may provide a more robust method for identifying the defect wall.

However, it is important to remember that we ultimately need to be able to measure the accuracy of our final mesh in order to properly quantify and report our findings. Having a robust method to do this, such as discussed in Aspert et al., would greatly increase the confidence we could have in our findings and also render them more communicable as we would have a standard measure of surface-to-surface error to report.

Next, I discuss the mathematical and algorithmic approaches described in the paper as well as their results and my assessment.

Mesh Simplification – Hoppe et al.

Background

Throughout the paper, the Hoppe et al. represents the mesh as simplical complexes and vertices (M = (K, V) where M is the mesh, K is the simplical complexes, and V is the vertices). Simplical complexes define the

surface of a mesh and are the space of unions between points, lines, and faces. Points are 0-simplices ($\{i\} \in K$), lines are 1-simplices ($\{i, j\} \in K$), and faces are 2-simplices ($\{i, j, k\} \in K$). We are now ready to begin discussion of the methods in Hoppe et al.

Mathematical Approach

To simplify a mesh, an energy function was minimized using legal moves over the mesh space. This was iterated until convergence. To reduce computation time, localization was used so that the energy function would not have to be evaluated of the whole mesh.

Energy Function

To simplify a mesh, Hoppe et al. minimizes the following energy function:

$$E(K, V) = E_{dist}(K, V) + E_{rep}(K) + E_{spring}(K, V)$$

Where E_{dist} is the sum of all squared distances from the points of the new mesh to the surface of the old mesh, $\sum_{i=1}^{n} d^2(\mathbf{x}_i, \phi_V(|K|))$. E_{rep} is the number of 0-simplices, or vertices, in the new mesh with a weight of c_{rep}, $c_{rep}m$ (where m is the number of vertices in the resulting mesh). E_{spring} is a concept introduced in this paper and is a spring force that pulls the vertices towards each other with a weight of k, $\sum_{\{j,k\} \in K} \kappa \|\mathbf{v}_j - \mathbf{v}_k\|^2$. E_{spring} is needed to ensure that there is always a local minimum for the energy function to achieve. Its inclusion prevents errors associated with an energy function that has no minimum which is illustrated below. The paper notes that the value of k can be reduced in later iterations after the mesh is in the neighborhood of its final value.

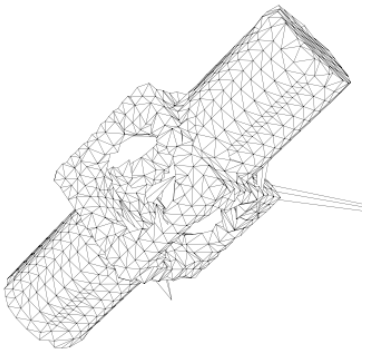

Example of a mesh without E_{spring} whose energy function did not have a minimum.

Legal Moves

The paper outlines three legal moves that can be applied to the mesh in order to achieve the minimum of the energy function. These are edge collapse, edge split, and edge swap. They are illustrated below with their respective formulas.

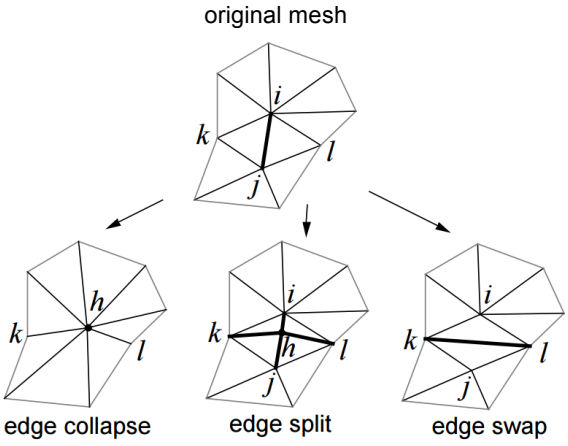

Edge collapse: $v_h = 0.5(v_j + v_i)$ with all edges that previously terminated in j or i now terminating in h.
Edge split: $v_h = 0.5(v_j + v_i)$ with i and j now sharing an edge with h and all vertices that share an edge with i and j now sharing an edge with h.
Edge swap: The edge between i and j is terminated and is replaced with an edge between k and l (vertices that are adjacent to i and j, on the same face, and were previously unconnected).

Localization

In theory, the energy of the entire mesh would be evaluated after each legal move, but computing the energy of the entire mesh is too inefficient to be practical. The paper uses the method of localization to reduce this problem. Though they did not go in depth into the mathematical method used, the basic concept is that a change a local area of a mesh will not affect vertices far away (see figure below).

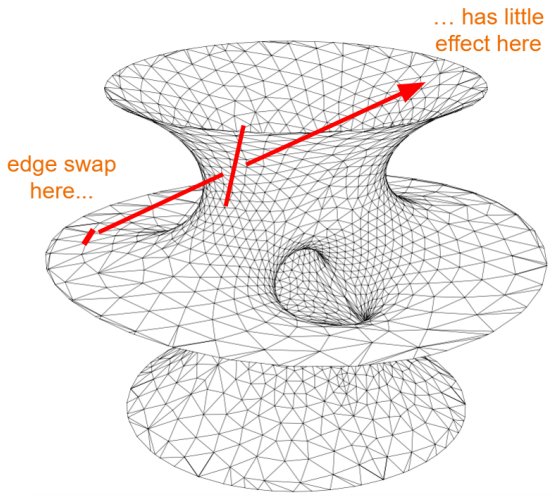

Thus, only the energy function a submesh of the local area of a legal move is evaluated to determine the which legal move to use to create the next iteration. This is much more efficient and yields results on par with evaluating over the whole mesh.

Algorithmic Approach

The algorithmic method used to implement the mathematical concepts is outlined below.

Pseudo Code

The following is the pseudo code ideal for implementing this mesh simplification process that is used in the paper:

```
OptimizeMesh(K0,V0) {
    K := K0
    V := OptimizeVertexPositions(K0,V0)
    – Solve the outer minimization problem.
    repeat {
        (K′,V′) := GenerateLegalMove(K,V)
        V′ = OptimizeVertexPositions(K′,V′)
        if E(K′, V′) < E(K, V) then
            (K,V) := (K′,V′)
        endif
    } until convergence
    return (K,V)
}

– Solve the inner optimization problem
–   E(K) = min_V E(K, V)
– for fixed simplicial complex K.
OptimizeVertexPositions(K,V) {
    repeat {
        – Compute barycentric coordinates by projection.
        B := ProjectPoints(K,V)
        – Minimize E(K, V, B) over V using conjugate gradients.
        V := ImproveVertexPositions(K,B)
    } until convergence
    return V
}

GenerateLegalMove(K,V) {
    Select a legal move K ⇒ K′.
    Locally modify V to obtain V′ appropriate for K′.
    return (K′,V′)
}
```

Idealized pseudo code for mesh simplification

First, the initial mesh is taken in and the vertex positions of the reconstructed mesh are adjusted to reach the minimum distance from the original mesh (this is the OptimizeVertexPositions function and is the "inner minimization problem"). Next, the program evaluates the possible legal moves and for each one evaluates the energy function after the vertex positions are optimized. The legal move that produces the lowest energy function is accepted as the next iteration of the optimized mesh and the process is repeated until convergence is reached (this is the “outer minimization problem”).

Results

The following is a table of results of this method applied to various mesh topologies and energy function values. A table of figures of the initial and final meshes for each table entry is also included.

Fig.	#vert. m	#faces	#data n	Parameters c_{rep}	Parameters κ	Resulting energies E_{dist}	Resulting energies E	time (min.)
7c	1572	3152	4102	-	-	8.57×10^{-2}	-	-
7e	1572	3152	4102	10^{-5}	10^{-2}	8.04×10^{-4}	4.84×10^{-2}	1.5
7f	508	1024	4102	10^{-5}	10^{-2}	6.84×10^{-4}	3.62×10^{-2}	(+3.0)
7g	270	548	4102	10^{-5}	10^{-3}	6.08×10^{-4}	6.94×10^{-3}	(+2.2)
7h	163	334	4102	10^{-5}	*varied*	4.86×10^{-4}	2.12×10^{-3}	17.0
7k	9220	18272	12745	-	-	6.41×10^{-2}	-	-
7l	690	1348	12745	10^{-5}	*varied*	4.23×10^{-3}	1.18×10^{-2}	47.0
7o	4059	8073	16864	-	-	2.20×10^{-2}	-	-
7p	262	515	16864	10^{-5}	*varied*	2.19×10^{-3}	4.95×10^{-3}	44.5
7q	2032	3832	-	-	-	-	-	-
7s	487	916	6752	10^{-5}	*varied*	1.86×10^{-3}	8.05×10^{-3}	9.9
7t	239	432	6752	10^{-4}	*varied*	9.19×10^{-3}	4.39×10^{-2}	10.2

Table of results

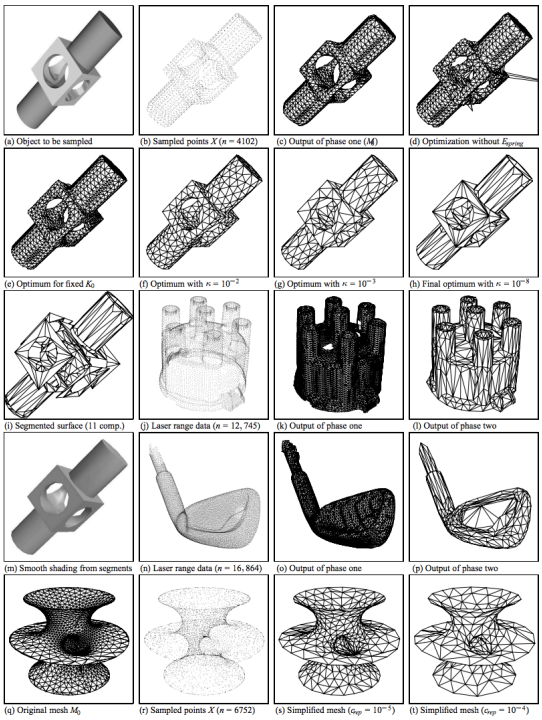

Table of figures

The resulting distances were on the order of magnitude of 10^-3 which may be sufficient for our application. Important to note is that a higher value of c_{rep} (which corresponds to the complexity of the resulting mesh) had a significant impact on the final distance energy, increasing it from 1.86 x 10^-3 to 9.19 x 10^-3.

Assessment

Overall I thought that this paper was well written and presented a versatile method for mesh simplification that is applicable to many applications. They broke the definition of their terms and methods into easily understandable sections and maintained a linear thought process through their explanation. However, there were also some points on which I believe this paper could be improved.

Some of the positive aspects of this paper were that their data representation (meshes and point clouds) was analogous to the data representation for our project which would make implementation of this method easier. In addition, this method allows for the recovery of sharp edges which is important for our ultimate goal of isolating the cranioplasty defect wall. Unlike some mesh simplification algorithms I looked into, this method allows for both the deletion and addition of points which allows for tighter control of overall distance from the original mesh. Similarly, the ability to weight the terms in the energy function allows for tighter control over the metrics of the final mesh. As our project is a medical application, this is important as we cannot readily sacrifice accuracy for simplicity.

Though for the most part this paper was efficient at conveying the information necessary to understand their method, they also included a lot of extrenuous concepts that weren't particularly necessary to understanding their method and only came up sparingly in their discussion such as the barycentric representation of elements. Also I did not feel that E_{spring}, one of the unique features of their method, was discussed in depth enough. They explained the necessity of having a third term that ensured a minimum, but why they used E_{spring} specifically was not touched on other than the kind of arbitrary reasoning that "it worked". I would have been interested in learning about the effects that using a different third term or having

additional energy function terms would have had on the overall results (neither of these ideas were discussed). Locality is also not discussed in depth despite being an important part of their algorithm. They describe it as only checking the energy function over a "submesh of the neighboring vertices", but the exact definition of what this means is not given. I also would have liked to see more noisy initial models for their test cases as this would have more closely matched what we are interested in using this method for. More generally, it is not unrealistic to assume that a starting mesh would be more irregular than the ones they have as representative of their method's capabilities. Finally, there was very little discussion on their mesh-to-mesh distance calculation. They use the sum of squared differences, but make no note on the potentially large effect using a different distance energy equation (such as max distance) could have on the final result. This is very important to our application as we need to find an accurate and robust gauge of distance. I supplemented this paper by reading another one that exclusively discusses error calculation (Aspert et al.).

Additionally, just as a separate note, it remains to be seen for our project whether the time it would take to simplify our meshes would outpace the time it would take to simply run our computations over an unsimplified mesh. If there is not a gain in the overall computational time, using it for that end is moot, although there is still promise in using it for segmentation.

Measure of Error – Aspert et al.

In this paper Aspert et al. begin by explaining the importance of being able accurately gauge the distance between meshes. They outlined some common methods of doing this (mean square error and total square error) but then suggested that a variation of the Hausdorff distance might provide a more accurate measure of error. Finally they outlined an efficient method for finding this distance. Note that this paper is meant to supplement the first paper and thus the discussion is not as in depth.

Background

The Hausdorff is the max Euclidean distance from surface to surface. Below are the equations from finding the distance from mesh S to mesh S'.

Euclidean norm: $d(p,\mathcal{S}') = \min_{p' \in \mathcal{S}'} \|p - p'\|_2$ Hausdorff distance: $d(\mathcal{S},\mathcal{S}') = \max_{p \in \mathcal{S}} d(p,\mathcal{S}')$

We are now ready to discuss the methods outlined in Aspert et al.

Mathematical Approach

Hausdorff Distance

As noted before the Hausdorff distance is $d(\mathcal{S},\mathcal{S}') = \max_{p \in \mathcal{S}} d(p,\mathcal{S}')$. However, it is important to note that this is not a symmetrical distance, (see figure below). In other words, $d(\mathcal{S},\mathcal{S}') \neq d(\mathcal{S}',\mathcal{S})$

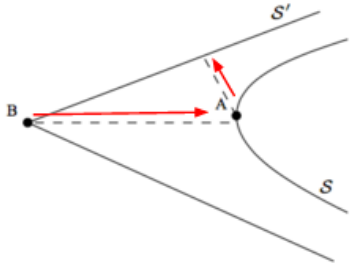

Example of non-symmetrical distance

This paper suggests that a rigorous method of distance calculation is to take the max of the two distances, max*[d(S, S'), d(S', S)]*.

Algorithmic Approach

Grid Sampling

In order to quickly compute the Hausdorff distance over a large mesh, the paper outlines a method for grid sampling. This allows for a discrete representation of the surface integral to compute the Hausdorff distance.

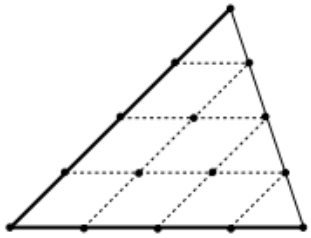

Example of n = 5 grid sampling

Results

Below are the results of the paper. They compared the forward and backward Hausdorff distances and noted that they could diverge significantly. This showed that taking the max of the two might provide more rigorous error estimation. Then they compared the time to compute the error to an analogous method that did not use the grid sampling. Their method (Mesh) was significantly faster.

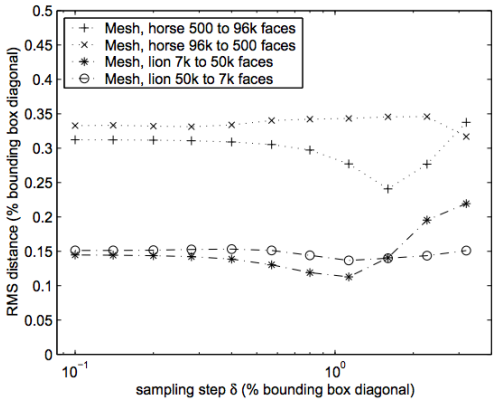

Differences in metrics of error

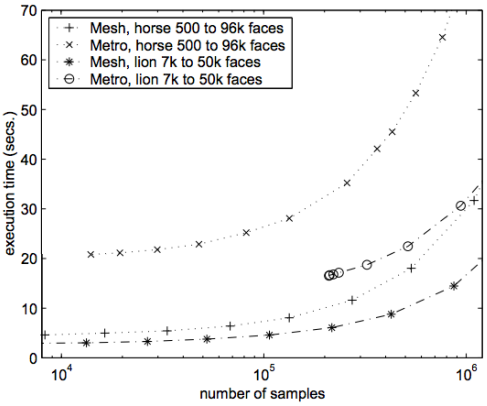

Time to compute error

Assessment

This paper did an excellent job of outlining a distance finding method that would work for a wide variety of topologies, not just more regular ones. It gave special attention to unusual corners and curves that could potentially be inconsistent when using other error finding methods. In addition it gives an "upper bound" on the error estimate that would be useful in our application where we want to limit the maximum error at all times. However, this method is more complicated than the wide spread max or mean error finding and its usefulness would be dependent on the application. Also, though the grid sampling greatly reduces computational time it cannot completely exhaust the surface meaning some inaccuracy would be possible.

Conclusion

The Hoppe et al. paper may provide a useful way to simplify our mesh and an alternative segmentation method. However, it is unknown whether our data would respond well to this optimization method or if the time tradeoff to compute the simpler mesh would be worth it in the long run. Similarly, the Aspert et al. provides a more rigorous method for distance calculation that we need to find the accuracy of our defect wall, but it remains to be seen whether it will provide a much different result for our particular kind of mesh. Both however provide interesting solutions to challenges in our project and we will be gauging the feasibility and utility of implementing them in the upcoming weeks.