CS749: Digital Geometry Processing

End-semester examination, Spring 2016

Question 1. (Surface Curvature)

(a) A large sphere (positive Gaussian curvature) has a small saddle-shaped patch (negative Gaussian curvature) smoothly embedded on its surface (see figure below). The point p is at the center of the saddle, where the dark lines cross. We draw a sequence of circles $C_1, C_2, C_3, \ldots$ on the surface, all centered at p, each slightly larger (in radius) than the last. C_1 has an infinitesimally small radius. What can you say about the sequence of areas of the circles, as compared to their radii? Similarly, what can you say about their circumferences?

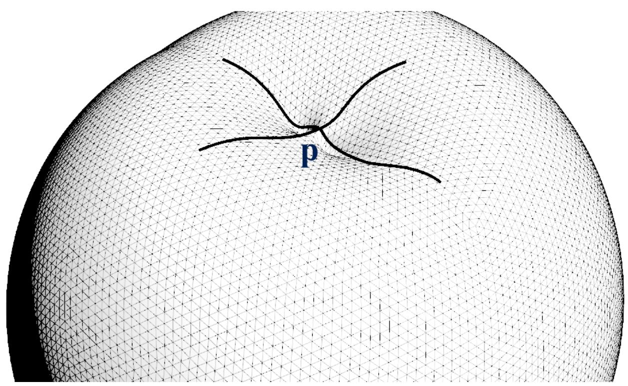

(b) A (closed) volume V in $\mathbb{R}^3$ is *convex* if every pair of points in it are connected by a straight line segment that lies fully within V. What can you say about the Gaussian curvature of the boundary surface ∂V of V? Justify your argument with pictures if necessary, formal proof not required.

(c) A surface point has principal curvatures κ_1 and κ_2, with $|\kappa_1|$ strictly greater than $|\kappa_2|$. Write the signs (positive, negative or zero) of its Gaussian and mean curvatures, if known, for each of the following situations.

(i) $\kappa_1 > 0$, $\kappa_2 > 0$

(ii) $\kappa_1 > 0$, $\kappa_2 < 0$

(iii) $\kappa_1 < 0$, $\kappa_2 > 0$

(iv) $\kappa_1 < 0$, $\kappa_2 < 0$

(d) A surface undergoes an isometric deformation. What happens to its first and second fundamental forms?

Question 2. (Point Clouds)

(a) A sphere S is approximated by a triangle mesh M with n vertices, where every vertex lies exactly on the surface of the sphere.

(i) Is it possible to construct a point cloud P_M, such that, *even without knowing the order of points* in the point cloud, it is possible to *exactly* reconstruct the geometry and topology of the mesh M? If yes, present such a construction (with as few points as possible) and argue why it is correct. If no, argue why not. Note again that you cannot exploit the order of points in P_M, i.e. you should assume that they are randomly shuffled before being passed to any proposed reconstruction algorithm.

(ii) If there was an additional constraint that the points in P_M must all lie on the surface of S, would your answer to part (i) change in any way? Justify.

(iii) If mesh M approximated an unknown, arbitrary 2D manifold, would your answer to part (i) change in any way? Justify.

For all parts of this question, assume that computations are numerically exact, so floating-point precision is not an issue.

(b) Orthogonal regression is one way of fitting a plane to a distribution of points. In class, we looked at a least squares solution: the optimal plane H_2^* is:

$$H_2^* = \underset{H \in \text{Planes}}{\arg\min} \sum_{i=1}^{n} \|\mathbf{x}_i - \mathbf{h}_i\|^2$$

where $\mathbf{x}_i$ is the i^{th} point, $\mathbf{h}_i = \mathbf{x}_i - (\mathbf{n} \cdot (\mathbf{x}_i - \mathbf{a}))\mathbf{n}$ is its closest point on the candidate plane $H = (\mathbf{n}, \mathbf{a})$, and $\|\mathbf{x}_i - \mathbf{h}_i\|$ is the Euclidean distance between these two points. Imagine, instead, that we had minimized a different error:

$$H_p^* = \underset{H \in \text{Planes}}{\arg\min} \sum_{i=1}^{n} (\|\mathbf{x}_i - \mathbf{h}_i\|_p)^p$$

where $\|\mathbf{u}\|_p$ is the L_p-norm of $\mathbf{u} = (u_1, u_2, \ldots, u_d)$, defined as

$$\|\mathbf{u}\|_p = (|u_1|^p + |u_2|^p + \cdots + |u_d|^p)^{\frac{1}{p}}$$

(i) Would you get the same optimal plane in all cases, i.e. $H_2^* = H_p^*$ for all p? Justify your answer – pictorial arguments are sufficient.

(ii) To minimize the error in the "p-case", your friend suggests that you compute the gradient of the error in the same way as we did for the sum of squares. For what values of p would this always work, and what values would it not?

(c) Researcher A has discovered a way to parametrize the space of human body shapes using 20 real numbers (height, weight, and so on) and further, has devised an efficient function F to uniquely map 20 arbitrary 3D points to the corresponding body shape containing all of them. Researcher B plans to employ RANSAC to detect humans in a large 3D scan, using function F to find the body shape supported by 20 randomly chosen scanned points and then checking if it is consistent with surrounding data. Is this a good idea? Why or why not?

Question 3. (Polygon Meshes)

(a) The geodesic distance between two vertices of a polygon mesh is approximated by the shortest path between the vertices in the graph of mesh edges. Is it possible to say if the approximation is always an upper bound or a lower bound on the geodesic distance? Briefly justify your answer.

(b) Sketch a 2D mesh in which the discrepancy between the shortest path distance and the true geodesic distance between two vertices is as large as possible. It should be possible to scale your construction to an arbitrary number of faces.

(c) Recall the definition of manifold mesh: A mesh is (topologically) manifold if

- Every edge is adjacent to 1 or 2 faces, and
- The faces around every vertex form a closed or open fan.

Is it possible to convert any non-manifold mesh to a manifold mesh, without changing the geometry of its individual polygons? If so, how would you do it?

(d) Sketch the shape corresponding to the following OBJ file. Assume Z is the up direction.

```
v  1   1   1
v  1  -1  -1
v -1   1  -1
v -1  -1   1
f 1 2 3
f 1 2 4
f 1 3 4
f 2 3 4
```

(e) A mesh, stored as a half-edge data structure, has k faces incident on vertex v. Write a loop (in C++-style pseudocode) that visits every face incident on the vertex in $O(k)$ time.

(f) A function f is defined on the surface of a shape M. You are given its spectrum, that is, the projections of f onto the eigenfunctions of the Laplace-Beltrami operator of M. Is it possible to exactly reconstruct the function?

(g) Two shapes are isometric to each other. What can you say about their Laplace-Beltrami eigenfunctions?

Question 4. (High-level shape analysis)

(a) You have a library that can compute the following global features on a shape:

- Shape histograms (SH)
- Histogram of pairwise geodesic distances (GD2)
- Lightfield Descriptors (LFD)

Which would be most natural to use for each of the following situations when searching a database of shapes? Justify each answer briefly.

(i) Find all human bodies, irrespective of pose.

(ii) Achieve the greatest accuracy, without needing to be pose-invariant or time-efficient.

(iii) Have the fastest preprocessing, since shapes are constantly being added to the database.

(b) Recall the supervised segmentation method using a conditional random field [Kalogerakis et al. 2010]. If the initial unary classification of faces was wildly incorrect, would the CRF optimization with pairwise terms still be able to compute the correct final segmentation/labeling? Why or why not?

(c) The parametrization methods studied in class require an open mesh, i.e. one with a boundary. How would you apply them to closed meshes, which are more commonly encountered in practice?

(d) Can the surface of the moon be perfectly represented with a heightfield (over a sphere)? Why or why not? Assume resolution is not an issue.

(e) Pixar has hired you to design their next-generation character creation tool. Describe 3 things that you would incorporate into this tool, that are significantly different from the way artists currently model characters.

(f) Apropos your previous answer, argue why Pixar might ***not*** want to change their current workflow.