Spring 2000
reading and research report

Stuart Andrews

May 15, 2000

Overview

This will be to the point.

Roughly, these were some of the goals that I proposed to accomplish at the beginning of the term (refering to Y2K-syllabus):

explicit goals

- Present Master's
- Extract break ridges from sherds using Master's research
- Propose a maximum likelihood, or graph theoretic solution to the sherd reconstruction problem
- Collaborate

implicit goals

- Learn about Jot and Java3D
- Develop interaction skills for collaborative research
- Learn about implicit surfaces

Some of these goals were completed. Some of them were not. The highlights of my semester certainly included my Master's presentation, and obtaining some preliminary results from the sherd CT data. On the other hand, the busy semester did not allow me to thoroughly explore the implicit surface literature, as I had planned. The specific goals that were reached are outlined below.

Accomplishments

Listed below are goals that were reached.

... what was done

- proposed the top-down framework for re-constructing pottery vessels
- presented my master's thesis
- extracted *break curves* and *etchings* from several sherds using interactive techniques
- segmented sherd surfaces into *structural components* using a min-cut segmentation algorithm
- collected and read a number of papers related to the pottery project (soon to be added to my bibtex file /pro/graphics/bib/stu.bib).
- read Bloomenthal's Implicit Surface book
- compiled a list of papers on implict surfaces that I would read
- collaborated (perhaps I need to be a little less flamboyant?)

... what was not done

- implementation of top-down framework for sherd reconstruction
- writing the first chapter of my thesis (although the title of first chapter is likely to be: 'The structure of surfaces', to be followed by a subsequent chapter: 'The structure of volumes'.)

More Details

... on break curves

To generate break curves on a sherd using 3D LiveWire techniques, we first had to convert the sherd data from the CT image into a useful surface format. Leymarie, Laidlaw and the software package AVS were instrumental in making this happen on short notice.

Once I had a triangle mesh surface representation, I could use my Master's code to extract the curves (and etchings). This involved:

1. computing an estimate of surface curvature
2. choosing a curve smoothness parameter
3. picking the endpoints of curve segments

Although these initial results provide us with sample data as well as valuable insight into the pot reconstruction problem, the fact is that a user must guide the process, and thus it can not be automated without further work. For this reason, I focussed my attention on another technique that accomplishes nearly the same thing.

... on segmentation

The goal of this experiment was to shoow that we could decompose a surface into structurally meaningful components. For example, a sherd may be divided into break surfaces, and a pair of inner and outer surfaces.

The break surface of a sherd is defined as the surface that is exposed when a piece of pottery breaks. For most of the sherd data, the break surface is usually a long and thin subset of the total sherd surface. Break surfaces may be used to measure how well two sherds fit together.

The inner and outer surfaces of a sherd refer to the smooth regions of the sherd's surface, which are named to correspondend with the inside/outside relationship of the vessel as a whole. These surfaces are surfaces of revolution, and hence they contain both a profile curve and a curve of revolution. These curves will be easy to estimate from the segmented inner and/or outer surface.

We would like some way of automatically segmenting the surface domain into distinct regions S and T ($S =$ break surface, and $T =$ inner $\cup$ outer surface). Other decompositions are also possible. As input, we are given a triangle mesh representation of the surface, and low-level measurements (curvatures) as the observed variables, y. To segment the surface, I have used a simple probabilistic model from statistical physics called the Ising Model (also used for 2D image segmentation).

The Ising model consists of two things:

- *a prior model* - $P(x)$ - which gives the likelihood of an arbitrary segmentation, $x = (S, T)$.
- *a data model* - $P(y|x)$ - which gives the likelihood of some observation y, given the segmentation is $x = (S, T)$.

Essentially, the prior $P(x)$ is defined to favour segmentations that have a short boundary (∂S) between the two regions. The data probability is defined to reflect how well the segmentation matches observed evidence y. We then pick the maximum a posteriori (MAP) estimate for the segmentation x. This is the maximum of $p(x|y) \propto p(y|x)p(x)$. The MAP estimate is the segmentation x that best fits the observed data, while respecting the prior of the Ising model. The solution is obtained using an equivalence between the segmentation problem and the problem of finding the Maxflow / Mincut in some derived network. The minimum cut determines the sets S and T by labeling the vertices with either +1 or -1.

On my sherd web page, I have several examples showing the curvature data used as the observed input vector y, and the corresponding segmentation $x = (S, T)$, where the sets are marked with one or two colours.