Progressive Shape Compression with Radial Basis Functions

Héctor M. Briceño, Thouis Jones, Leonard McMillan

Introduction:

Compression of geometric models is becoming more and more necessary as their complexity grows. It is not unusual for models to have millions, or even billions, of points [1]. To work with such models it is desirable, if not necessary, to employ some form of compression. Progressive compression is preferrable, such that an approximate shape can be produced given a few bytes and refined as more data is available.

We have developed a new progressive compression method for geometric models based on three key elements:

- Representing shapes by their signed distance fields;
- Using Radial Basis Functions for reconstruction of distance fields from unstructured samples;
- Stochastically drawing these samples from biased distributions to reduce the number required to represent a shape.

Our main contributions are an empirical determination of how to bias the sampling for effective compression of the samples, and an efficient progressive encoder combining the elements above.

Approach:

There are strong connections between distance fields and Voronoi methods, and we can use the criterion from [2] to develop a heuristic for placing samples of the distance field for optimized coding efficiency. The medial axis of a shape occurs at C^1 discontinuities in the shape's distance field. In our continuous RBF reconstruction, these discontinuities cause areas of higher curvature. We also want to concentrate samples near the shape's surface, since we are most concerned with the reconstructed surface's fidelity. To meet these two goals, we use

$$P(\mathbf{x}) = \frac{C(\mathbf{x})^\alpha}{|D(\mathbf{x})|^\beta + \epsilon} \tag{1}$$

as a weighting function for sample locations, where $C(\mathbf{x})$ is the interpolant's curvature at location $\mathbf{x}$, $D(\mathbf{x})$ is the signed distance function, and α, β are free parameters which we must optimize for our algorithm. ϵ is a small value present to prevent points on the surface from receiving infinite weight.

We combine distance fields, RBFs, and biased sampling into a compression method as follows. We start with some shape, represented as its signed distance function D. Our goal is to produce a set of sample locations and distance field values that can be represented with a minimum number of bits and such that the shape reconstructed from those samples is as accurate as possible.

At each step of the algorithm, we assume we have already generated a set S_n of samples:

$$S_n = \{(\mathbf{x}_i, D(\mathbf{x}_i))\}, \qquad i = 1, ..., n \tag{2}$$

At step $n+1$ of the algorithm, we improve our reconstruction by adding a new sample/distance pair $(\mathbf{x}_{n+1}, D(\mathbf{x}_{n+1}))$ to S_n to form S_{n+1}. The location of this sample is drawn from the biased distribution defined by equation 1, *calculated based on* S_n.

Thus, starting from some initial set S_k used to "seed" the RBF with a valid solution, we produce nested sets of samples and distances

$$S_k \subset S_{k+1} \subset S_{k+2} \subset ... \subset S_N \tag{3}$$

Drawing samples directly from the continuous biased distribution described by equation 1 would be quite difficult, as it is neither normalized nor easily invertible. Rather, at each step we form a discrete approximation by drawing a large number of uniform samples in the bounding box of the shape, weight them according to equation 1, normalize, and draw a sample from the resulting discrete distribution.

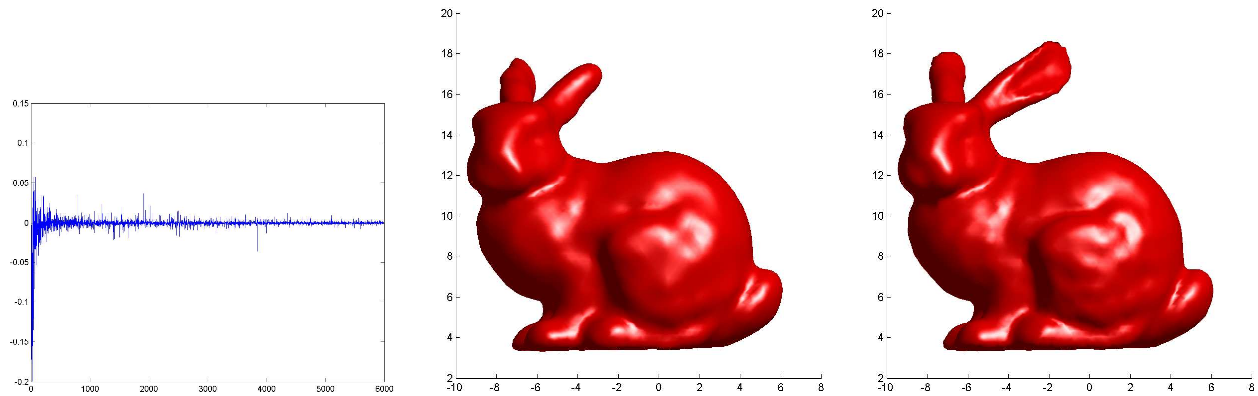

Figure 1: Left to right: error coefficients get smaller as our prediction model improves, and the Stanford Bunny reconstructed using 2000 and 6000 samples.

A sample location $\mathbf{x}_{n+1}$ chosen from this distribution is determined solely by the approximated shape at the previous step, as reconstructed from S_n (assuming a deterministic source of random numbers, i.e., a PRNG). Therefore, we can remove the locations $\mathbf{x}_{i>k}$ from S_n and represent it instead as the initial set S_k and the distances of the samples added at each step

$$S_n = \{S_k, D(\mathbf{x}_i)\}, \qquad i = k+1, ..., n. \tag{4}$$

These distances are not independent, but at each step $D(\mathbf{x}_{n+1})$ is predicted by the reconstruction from S_n. As n grows, the difference between the predicted and true distance values at $\mathbf{x}_{n+1}$ tends to grow smaller, as shown in Figure 1. These differences have much lower entropy than the distances themselves. We can quantize them to E_d (or our desired reconstruction accuracy) and compress them with a progressive arithmetic coder. Figure 1 shows the progressive nature of our encoding.

Progress: We have experimentally verified that biasing with equation 1 produces samples that accurately reconstruct detailed models. We have also empirically determined values for α and β. Our initial results are competitive with the best methods of shape compression, measured in rate-distortion.

Future: Our plan for future work is to optimize our code to accelerate compression and decompression. Achieving interactive rates will probably require adjustments to our algorithm above, and possible tradeoffs in compressibility.

References:

[1] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk, "The digital michelangelo project: 3d scanning of large statues," in *Proceedings of SIGGRAPH 2000*. July 2000, Computer Graphics Proceedings, Annual Conference Series, pp. 131–144, ACM Press / ACM SIGGRAPH / Addison Wesley Longman, ISBN 1-58113-208-5.

[2] N. Amenta, M. Bern, and M. Kamvysselis, "A new voronoi-based surface reconstruction algorithm," in *Proceedings of SIGGRAPH 98*, Orlando, Florida, July 1998, Computer Graphics Proceedings, Annual Conference Series, pp. 415–422, ACM SIGGRAPH / Addison Wesley, ISBN 0-89791-999-8.