Voxels And Stuff

[*Fast Multiresolution Image Querying*,
Charles Jacobs, Adam Finkelstein,
David Salesin]

Multiresolutional Analysis

- Describe the multiresolution approach for images
- Generalize this approach to three dimensional voxel grids
- Discuss

Multiresolution analysis for Images

- Generation of image signature
- Defining the image querying metric
- Specifying a data-structure for queries

Image Signature 1

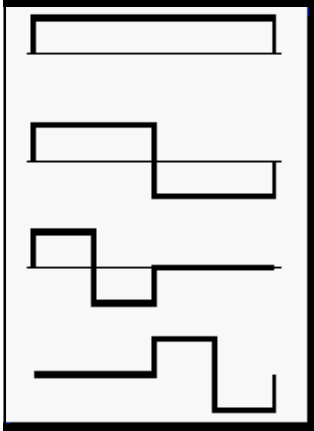

- Given an *n*x*n* image, generate an *n*x*n* array of the wavelet coefficients for the standard Haar basis functions. (Very fast to compute)

[http://www.cl.cam.ac.uk/~jeg24/PUBLICATIONS/SKETCHES/WAVELET/sld001.htm]

Image Signature 2

- Truncate: Find the *m* largest coefficients and set all others equal to zero
- Quantize: Set the non-zero coefficients to +1 or –1 depending on their sign

Jackie Chan Example

Original Image (256x256)

Truncated And Quantized to 5000

Truncated And Quantized to 1000

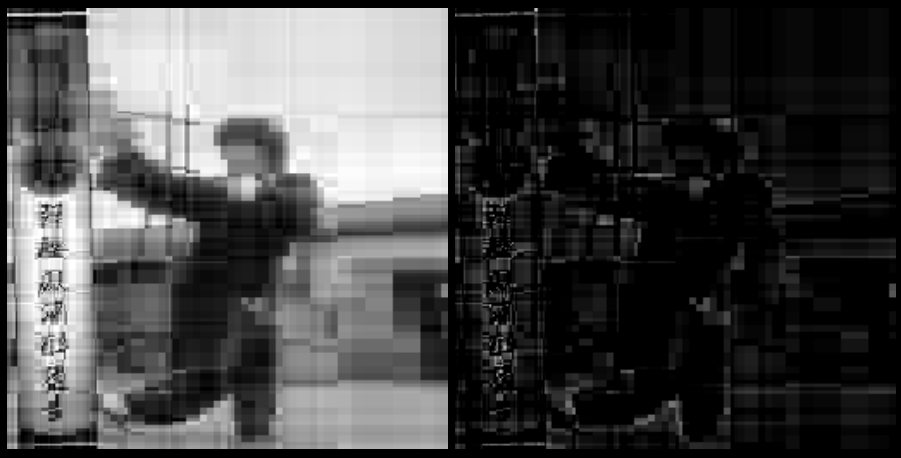

Truncated And Quantized to 500

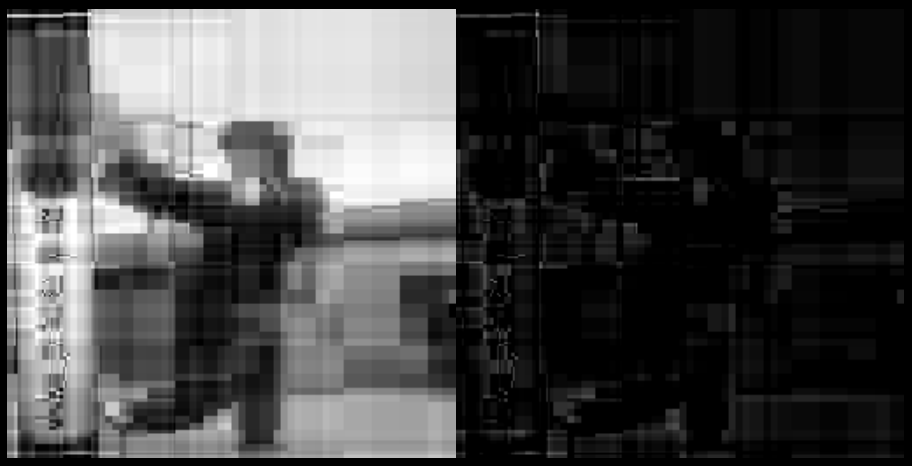

Truncated 100

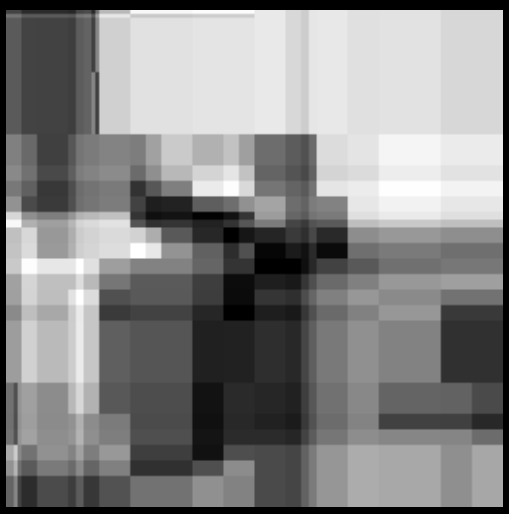

Truncated 50

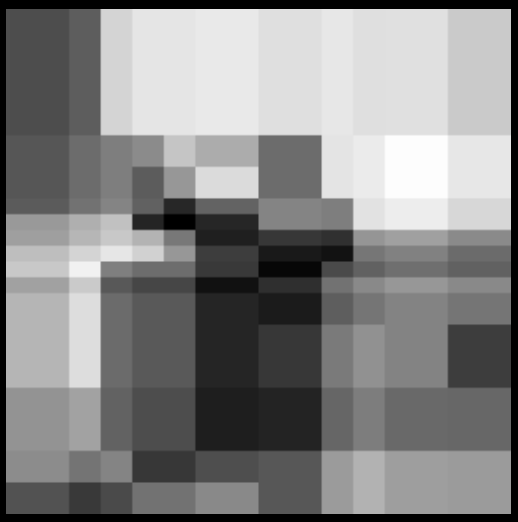

Truncated 10

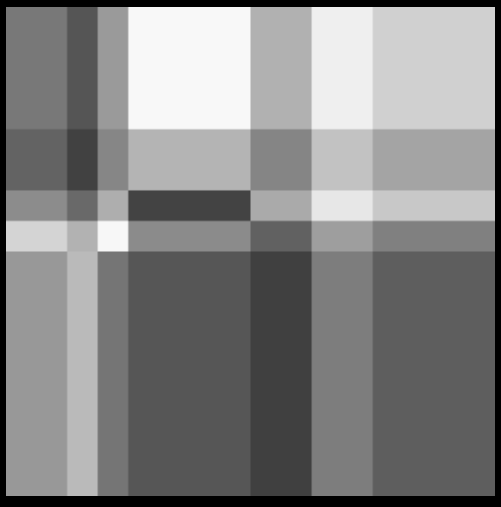

Querying Metric 1

The query metric is defined by:

$$\|Q,T\| = \sum_{i,j} w_{i,j} |Q[i,j] - T[i,j]|$$

where *Q[i,j]* and *T[i,j]* are the truncated and quantized coefficients and w_{ij} are weights, fine tuned to the database.

Querying Metric 2

It can be simplified to:

$$\|Q,T\| = \sum_{i,j:Q[i,j]\neq 0} w_{bin(i,j)} (Q[i,j] \neq T[i,j])$$

where *bin(i,j)* is a simple bucketing function dependent on the scale of the wavelet function to which they correspond

Data Structure 1

Preprocessing:

- The images in the database are truncated and quantized.
- Two 2-D arrays, D^+ and D^- are generated, with $D^+[i,j]$, respectively $D^-[i,j]$ indexing the list of images with high positive, respectively negative, wavelet coefficients.

Data Structure 2

Given a query image:

- The image is truncated and quantized giving a 2-D array Q with (-1,0,1) as entries
- A scoring array indexing all database elements is generated.
- For each indexing pair (i,j) with Q[i,j]>0, the elements in D^+ are used to update the scoring array (same for Q[i,j]<0)

Data Structure 3

The n best scoring database images are selected.

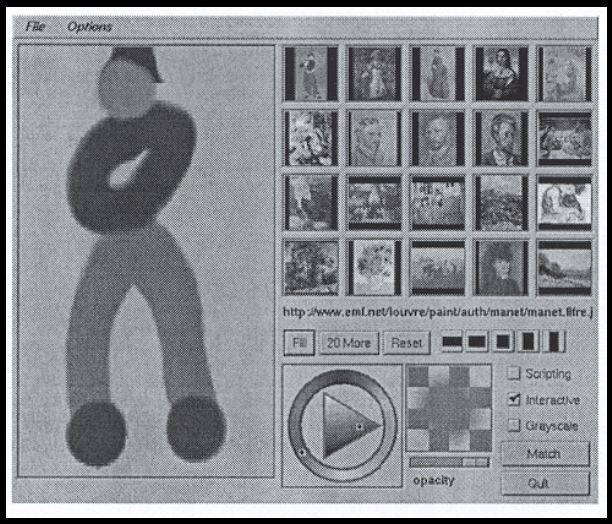

3D Generalization

The generalization of this method to a voxel grid is immediate. The big trick is establishing a good choice of weights.

[http://www.cl.cam.ac.uk/~jeg24/PUBLICATIONS/SKETCHES/WAVELET/sld001.htm]

The Magnificent Torus

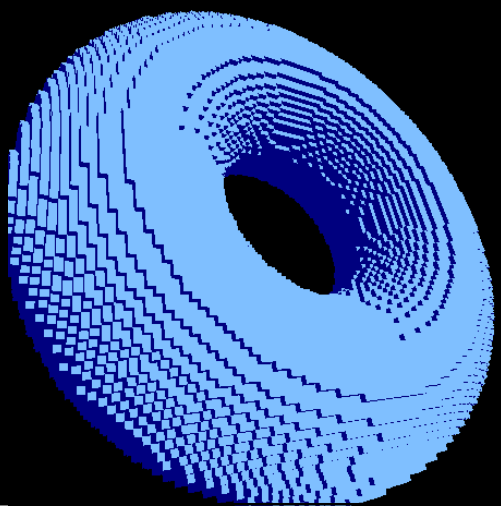

Torus Truncated to 5000

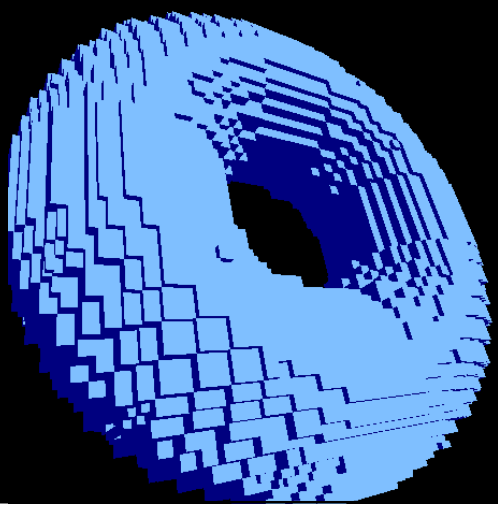

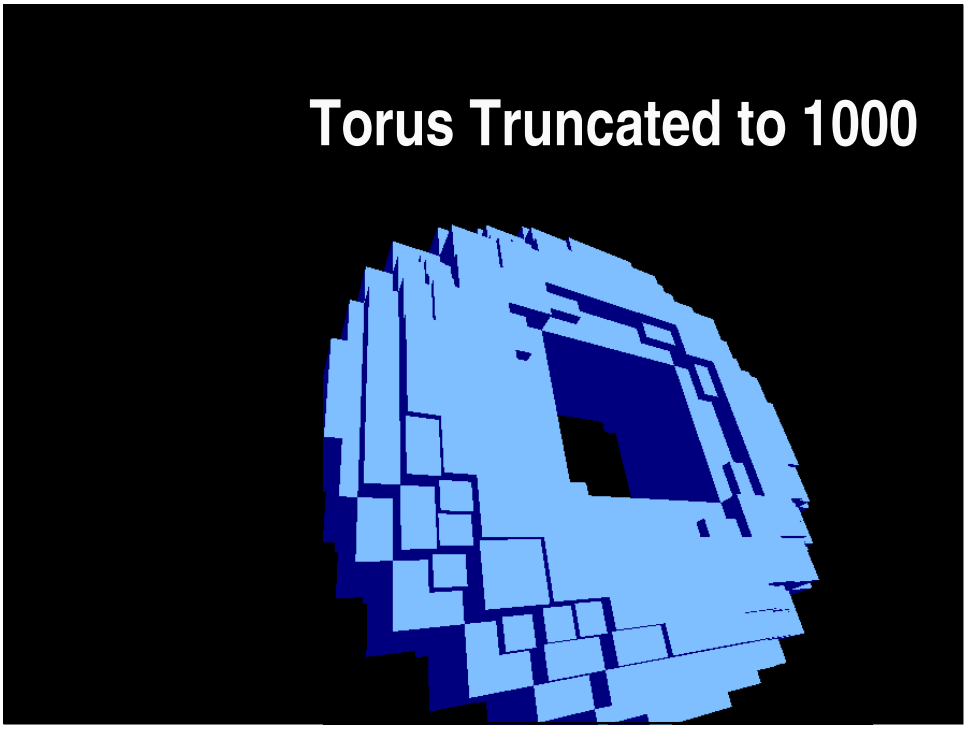
Torus Truncated to 1000

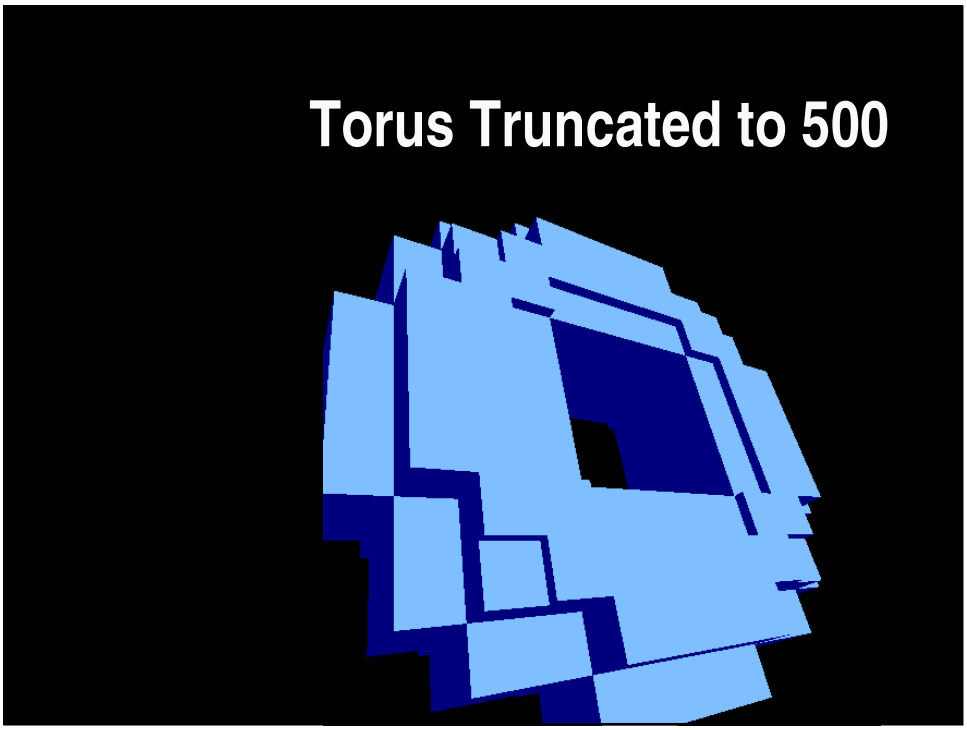
Torus Truncated to 500

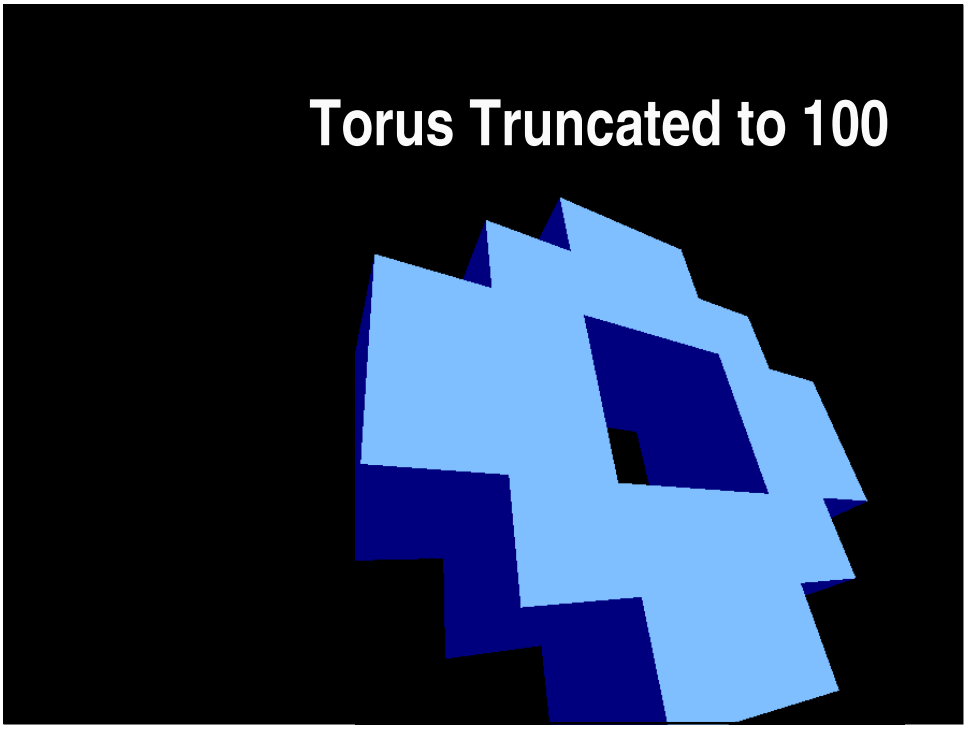
Torus Truncated to 100

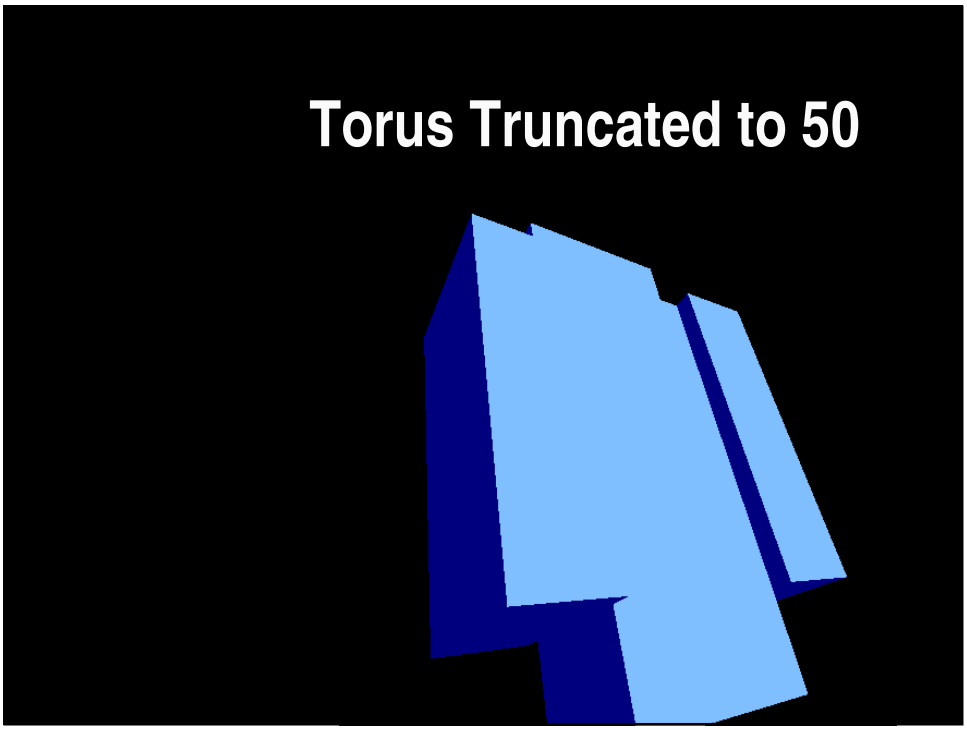
Torus Truncated to 50

Discussion 1

- Very Fast: (Works in a fraction of a second on databases of 20,000 images)
- The use of a Haar basis makes obtaining the signature very fast
- Invariant under small amounts of noise and perturbations

Discussion 2

- The query method is not hierarchical (i.e. O(n)) and hence is not satisfactory for large image/voxel databases (e.g. the web)
- It does not allow for affine transformations
- The Haar basis is anistropic

Discussion 3

- Even with a guarantee that it finds roughly the true target within 1% of the database, this becomes ineffective for large databases.
- The weights for the “metric” are determined after coefficients are discarded.