Lighting Transfer Functions Using Gradient Aligned Sampling

Eric B. Lum & Kwan-Liu Ma

September 20, 2006

"Most research in the use of higher dimensional transfer functions has focused on the specification of color and opacity, despite the fact that transfer functions can be used for specifying any of a number of optical properties including surface illumination."

- User specified surface illumination using a higher-dimensional transfer function, with input that consists of samples read along the gradient direction.
- An interface for specifying lighting transfer functions.
- All this to allow user to position lighting at surfaces independent of surface opacity.
  - Basically, this is not only for transitions between opaque and transparent materials, but also for transitions between two similarly transparent materials.

Post-Filtered

- Normals used for lighting are computed from the gradient of scalar values after applying the density lookup.
- Has the advantage of more closely following the physical interaction of light in a volume density.
- Requires the recomputing of gradients when the classification function is changed.

Pre-Filtered

- Yields normal directions that are more closely tied to the original data.
- Not influenced by potential errors in classification.
- Can be generated in a preprocessing step.
- Do not need to be recalculated if classification function is changed.
- Yields normals thar are oriented in directions contrary to densities during rendering.
- This paper presents a hybrid approach.

Sampling

- The transfer function takes into account data values perpendicular to a possible surface at a sample.
- This is done by sampling points along a sample's normalized gradient direction.
- Two additional scalar values are used: one in the direction of the gradient and one in the opposite direction.
- The center sample is used for coloring, and the additional samples for lighting.

Coloring and Lighting

- A transfer function colormap is used along with the Phong shading model.
- The usual lighting coefficients are replaced lookup tables that are functions of the two additional samples.
- The color of a rendered sample is then: $C = ColorTF(S)(LTF_{k_a}(S_1, S_2) + LTF_{k_d}(S_1, S_2)MAX(N \cdot L, 0)) + LTF_{k_s}(S_1, S_2)MAX((N \cdot R)^n, 0)$

(Cont.)

- The two samples read along the gradient direction are well-suited inputs for the lighting transfer functions:
  - If the two samples have the same value, the center sample is likely in a homogeneous region.
  - If the two samples belong to two different materials, the sample is likely on the boundary.
- The distance between sample pairs can be adjusted – useful if boundaries are thick.
- For datasets in this paper, the distance is always one voxel in each direction.

Transfer Function Histograms

- A histogram of a transfer function's domain can make specifying a transfer function easier.
- Plotting the scalar value pairs leads to an unintuitive histogram.
- Line primitives instead of point were used.
- Two parallel axes were used for the two scalar values.
  - The bottom axis corresponds to the scalar read in the gradient direction, and increases from left to right.
  - The top corresponds to scalar read in opposite direction.

(Cont.)

- Dense vertical lines represent homogeneous regions.
- Diagonal lines illustrate the boundaries between two materials.

Specifying the Transfer Function

- Done by assigning lighting coefficients to the different lines in the histogram.
- Pairs of 1D transfer functions are used for the two axes.
- The user specifies a reflectivity distribution for both the source of the lines at the top, and the destination of the lines at the bottom.
- These pairs are multiplied together to create separable 2D transfer functions that are summed to create the final 2D transfer function (not separable).

More on the Hybrid Approach

- User can specify transfer function that attenuates using either measure of homogeneity (post- or pre-classified).
- By not selecting vertical lines in the lighting transfer function, the user can attenuate surface lighting of homogeneous material in data space.
- By not specifying lines that link regions with similar opacity, attenuation of lighting based on post classifcation can be performed.

- Uses view-aligned textured polygons and a fragment program.
- Each sample's scalar value and gradient are read.
- The gradient is normalized, and the pair of additional scalr values along the gradient are read.
- The center value is used for color and opacity lookup.
- The scalar value pair are used as texture indices into the 2D lighting transfer function.

More Still on the Hybrid Approach

- The gradients from the original scalar field are precomputed.
- Based on the sign of the difference between the post-classified and sample densities, the gradient is adaptively flipped to point in the direction of maximum variation in post-classified density.
- This allows for single sided lighting to be used, getting features equivalent to two-sided ligthing without showing any back-facing surfaces.

- There are approximately twice as many texture reads required compared to traditional opacity-based transfer function methods.
- On Geforce 6800, this method gets 7.3 frames per second versus 13.5 frames per second with uniform lighting.