

VOLUME PADDING METHODS FOR VIEWING DATA AT ACQUISITION BOUNDARIES

When rendering a volume using VP500 hardware, the first slice or two at either end of the volume remains invisible. This is due to the fact that samples are visible only when they have a valid gradient. Since samples at the boundary do not have valid gradients, they will not appear in the result of the render. This could present a problem, for instance in applications where it is critical to visualize data at acquisition boundaries.

The only way to make all acquired data visible is to pad the volume with artificial slices. However, padding a volume can introduce unwanted visual artifacts in the rendered image. Thus, the recommended padding methods are augmented to eliminate such artifacts. The remainder of this document reviews the advantages, disadvantages, and implementation details of the recommended padding methods.

1) Data Padding + Turning off Lighting

Pros

All the data is displayed using this approach, and there are no image artifacts generated as a byproduct.

Cons The rendered scene is not lit.

Technical Implementation

The volume should be padded with three slices containing all zeros in each direction where the complete acquired data should appear. In addition, lighting should be disabled using the VLI API.

2) Data Padding + Cropping

Pros

All the data is displayed using this approach, and there are no image artifacts generated as a byproduct.

Cons Since cropping is used to frame the full volume properly, the cropping functionality becomes limited to box cropping only. Furthermore, crop boundary computations have to be adjusted to crop out the padded slices in addition to the desired data slices.

Technical Implementation

The volume should be padded with three slices containing all zeros in each direction where the complete acquired data should appear. In addition, cropping parameters should be set to crop out the extra slices.

3) Data Replication

Pros

All the data is displayed using this approach, and there are no image artifacts generated as a byproduct.

Cons

The size of the volume is increased by two slices, which adds time (minimal, but some) to rendering, locking and all other operations that are a function of volume size.

Technical Implementation

A total of two slices should be added to the data, one preceding the data, and one after the last slice of data. The very first slice of the new volume (also the first slice of padding) should be a copy of the second slice of data, and the very last slice of the new volume (also the second slice of padding) should be a copy of the second to last slice of data.

Given a volume dataset that has 'N' slices of data, where the slices are numbered 0, 1, 2, ... N-1 , :

1. load a single slice of pad data at the front of the volume that is a duplicate of slice #1.
2. load slices #0 through #N-1
3. load a single slice of pad data that is a duplicate of slice #N-2

This has the effect of minimizing the gradients on the first and last slice of actual data and thus the lighting film is diminished.

The order of slice data should look as follows:

slice #1 (pad)
slice #0
slice #1
slice #2
slice #3
...
slice #N-3
slice #N-2
slice #N-1
slice #N-2 (pad)

The volume is increased by a total of 2 slices.