Volume Rendering Volumetric Presentation State Supplement 190

DICOM Working Group 11

Letter Ballot

March 24, 2016

Volumetric Presentation States

- What is a Volumetric Presentation State (VPS)?
  - Standard representation of the 'recipe' to use in creating a 2D view of volumetric imaging data
  - Typical views of volumetric imaging data:
    - Multi Planar Reconstruction: View of a 'slice' through a volumetric dataset
    - Volume Rendering: View through volume as seen from a 'viewpoint', like a landscape view of the data

Volumetric View Examples

- Planar MPR:

- Volume Rendering:

Volumetric Presentation State Scope

- One VPS defines only one view and geometry
- Separate DICOM objects or ad hoc conventions are used to define layout of multiple views
  - Multiple VPS objects required to describe multiple views on one display screen (e.g., 3 orthogonal MPRs and one volume rendering).
  - Hanging Protocol or Structured Display object may formalize the layout of multiple views on one display
  - A hint is provided by a "display collection UID" in each related VPS
- One VPS may define animation from the initial view geometry

Presentation States vs. Layout

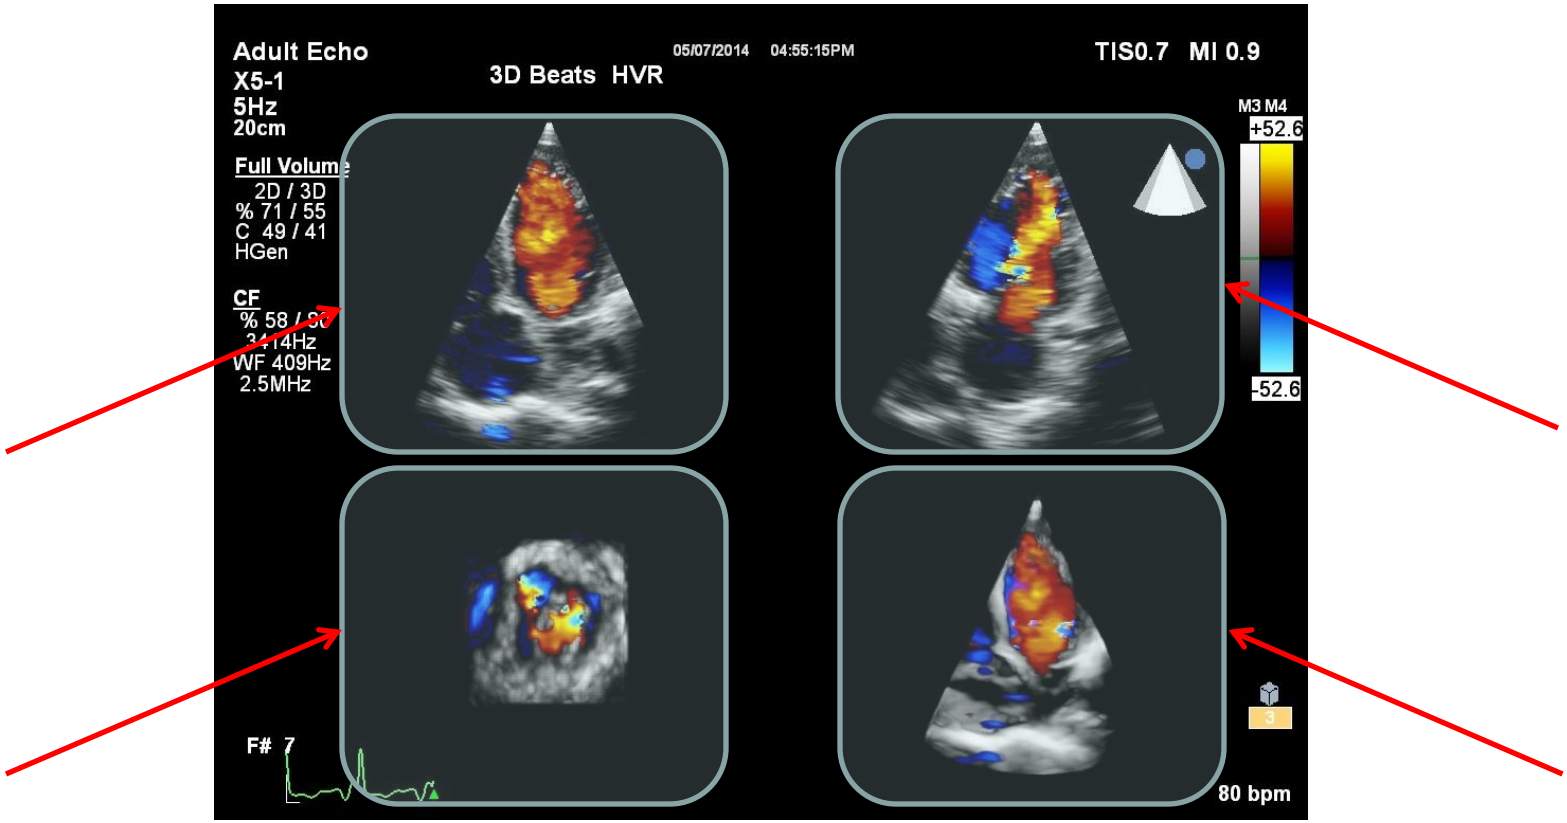

Each MPR or Volume Render view is created by One Presentation State
Layout of views on a display use a different mechanism.

Exact Match vs. Reasonable Consistency

- Application of a Volumetric Presentation State is not expected to be exactly reproducible on different systems, and an exact match of volume presentation on multiple devices cannot be guaranteed
- Reasonable consistency is provided by specification of inputs, geometric descriptions of spatial views, type of processing to be used, color mapping and blending, input fusion, and many generic rendering parameters, producing what is expected to be a clinically acceptable result.

Supplement 156 Planar MPR VPS

- VPS defines a Planar MPR view of volume data, with the following features:
  - Specification of planar geometry
  - Thin or Slab MPR
  - Cropping of unwanted volume data from the view
  - Blending and colorization of multiple volumes, including registration into a VPS display space
  - Clinical identification of the specified view
  - Text/graphic annotation in volume space or on 2D view
  - Animation of the volume from initial view
- Approved for Final Text in May, 2015

Supplement 190
Volume Rendering VPS

- VPS defines a Volume Render view of volume data with the following features:
  - Specification of the render geometry and lighting
  - Cropping of unwanted volume data from the view
  - Colored segmentation overlays of a single volume
  - Blending and colorization of multiple volumes, including registration into a VPS display space
  - Clinical identification of the specified view
  - Text/graphic annotation in volume space or on 2D view
  - Animation of the volume from initial view
- In Public Comment until early March, 2016

Standardization Challenges

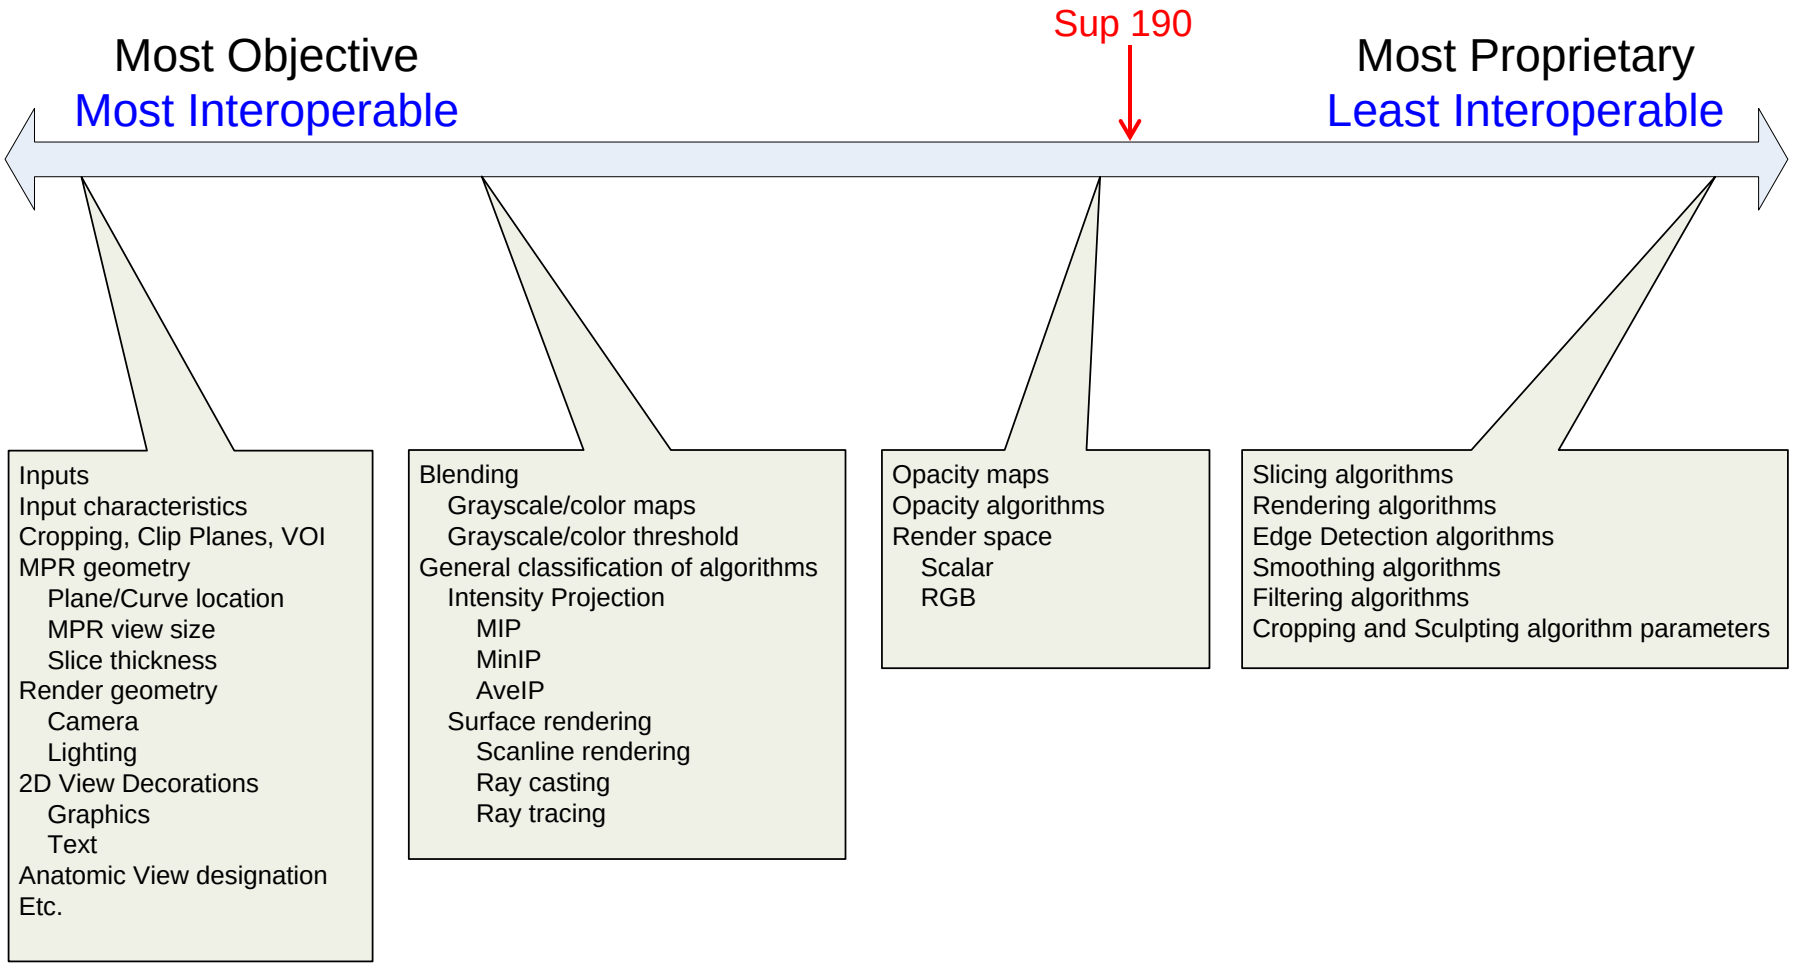

Sup 190 Volume Rendering SOP Classes

- Specifies three new VPS SOP Classes:
  - Volume Rendering VPS SOP Class
    - Single volume input
    - Bounding Box/Oblique Plane cropping only
    - Single display transformation
  - Segmented Volume Rendering VPS SOP Class
    - Single volume input
    - Multiple segmentations allowed with separate display transformations
  - Multiple Volume Rendering VPS SOP Class
    - Multiple volume inputs
    - Multiple segmentations allowed for each input volume input, with separate display transformations

Volume Rendering SOP Class

- Most basic SOP Class
- Used primarily for grayscale or colorized volume render view of single volume
- Global cropping allowed
  - Bounding Box
  - Oblique Crop Planes
- No per-input cropping – excludes Segmentation

Example for Pipeline

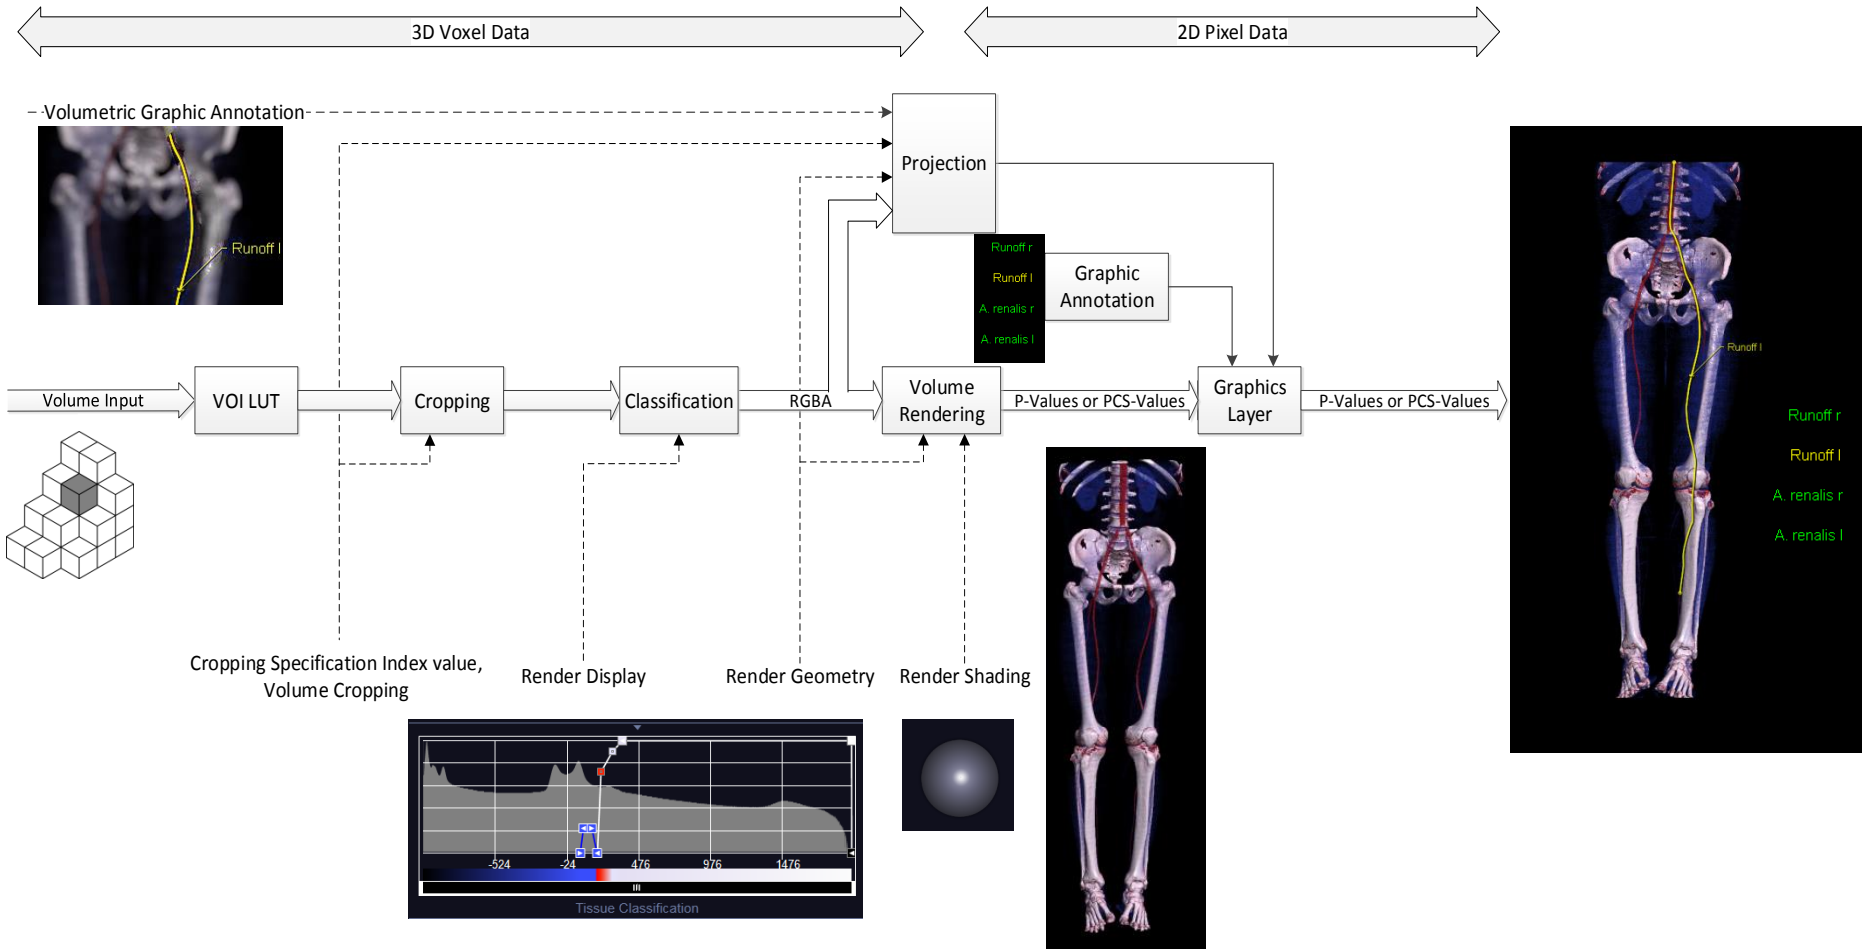

Segmented Volume Rendering SOP Class

- Intermediate complexity SOP Class
- Used primarily for grayscale render view of a single volume with one or more colorized segmentations
- All cropping styles allowed
  - Global Bounding Box and/or Oblique Plane cropping
  - Per-Input croppng of any type, including Segmentation

Segmentation Pipeline

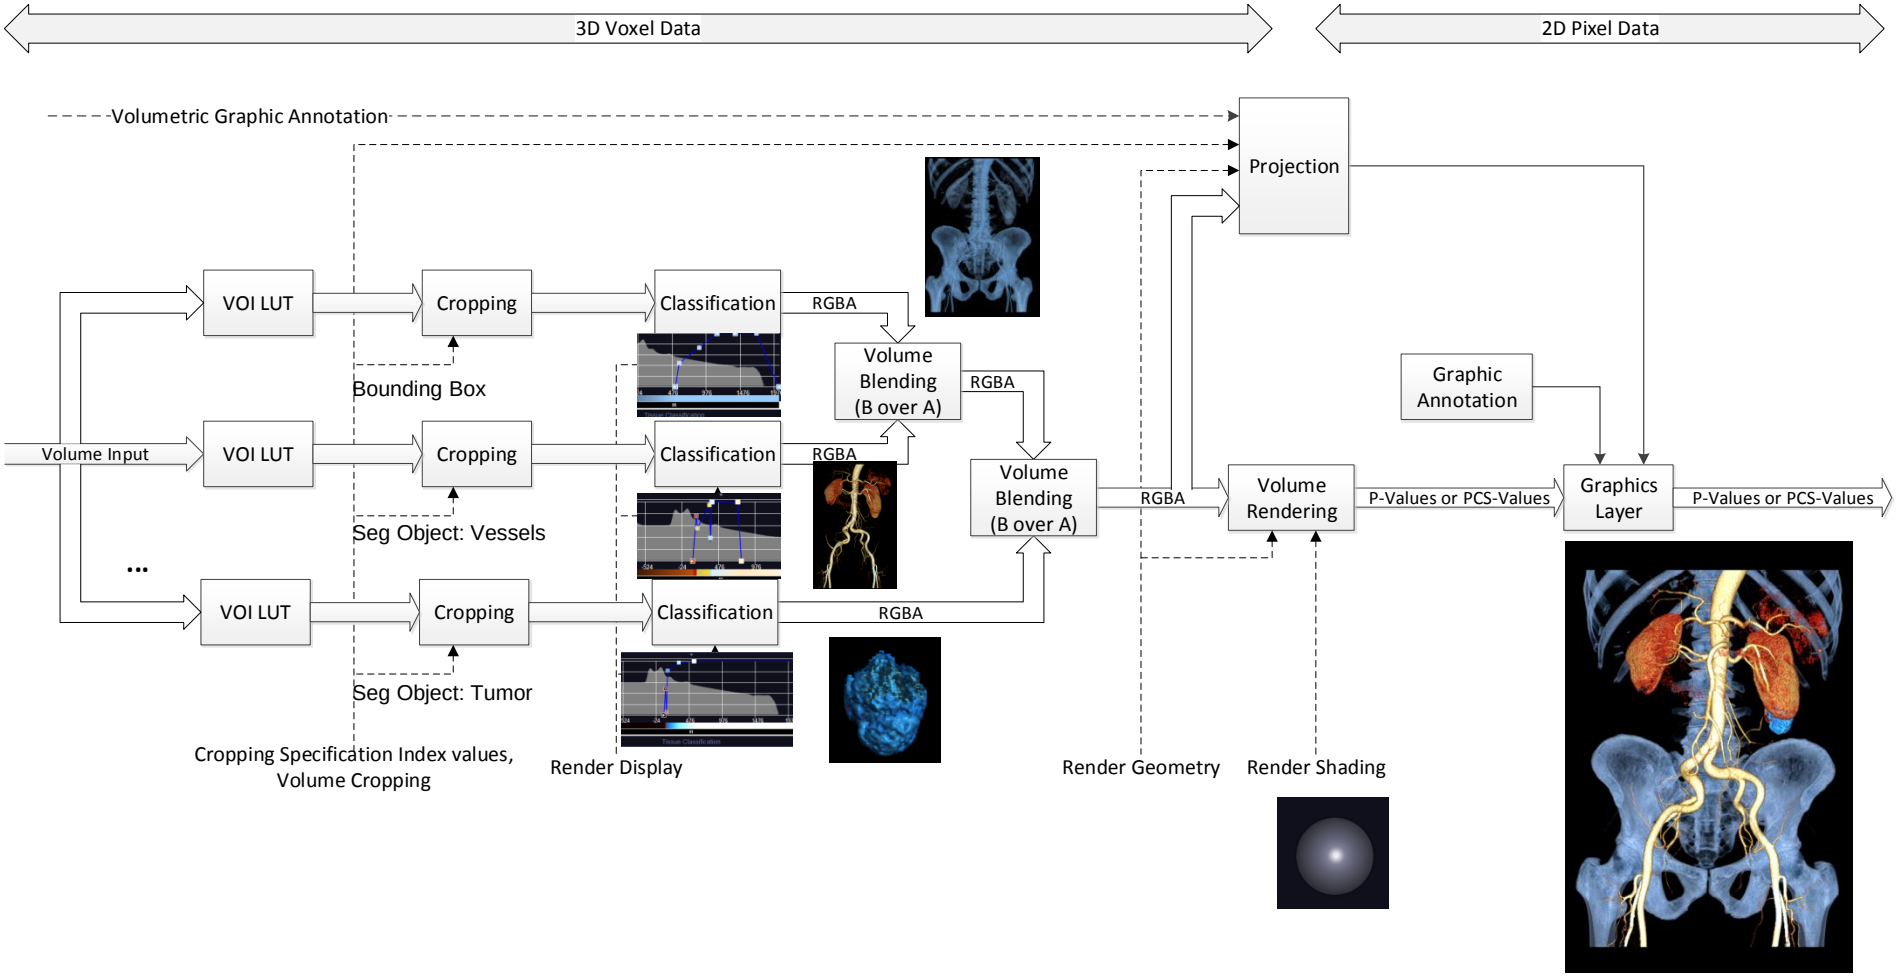

Segmentation Objects

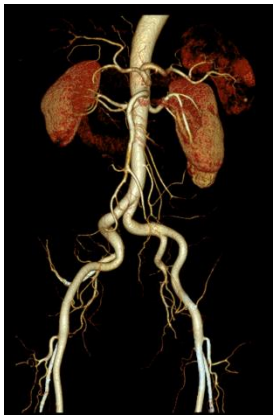

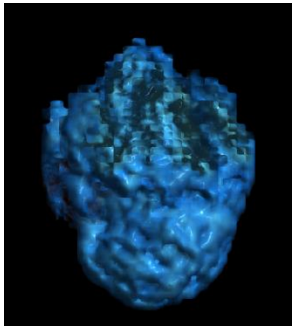

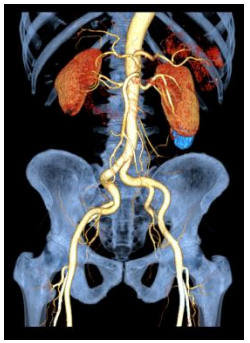

Anterior

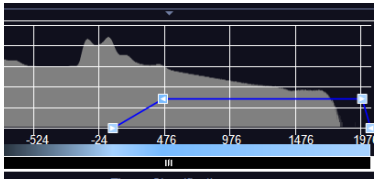

Bounding Box

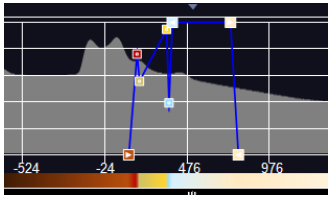

Seg Object: Vessels

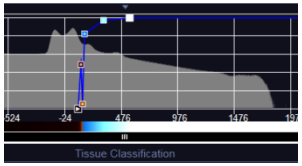

Seg Object: Tumor

Posterior

Multiple Volume Rendering SOP Class

- Most advanced SOP Class
- Used for specifying a volume render view of two or more “fused” volume input sets of the same or different modality
- Multiple “Volume Streams” define inputs to undergo RGBA Compositing
- All cropping styles allowed
  - Global Bounding Box and/or Oblique Plane cropping
  - Per-Input croppng of any type, including Segmentation

Multiple Volume Rendering SOP Class

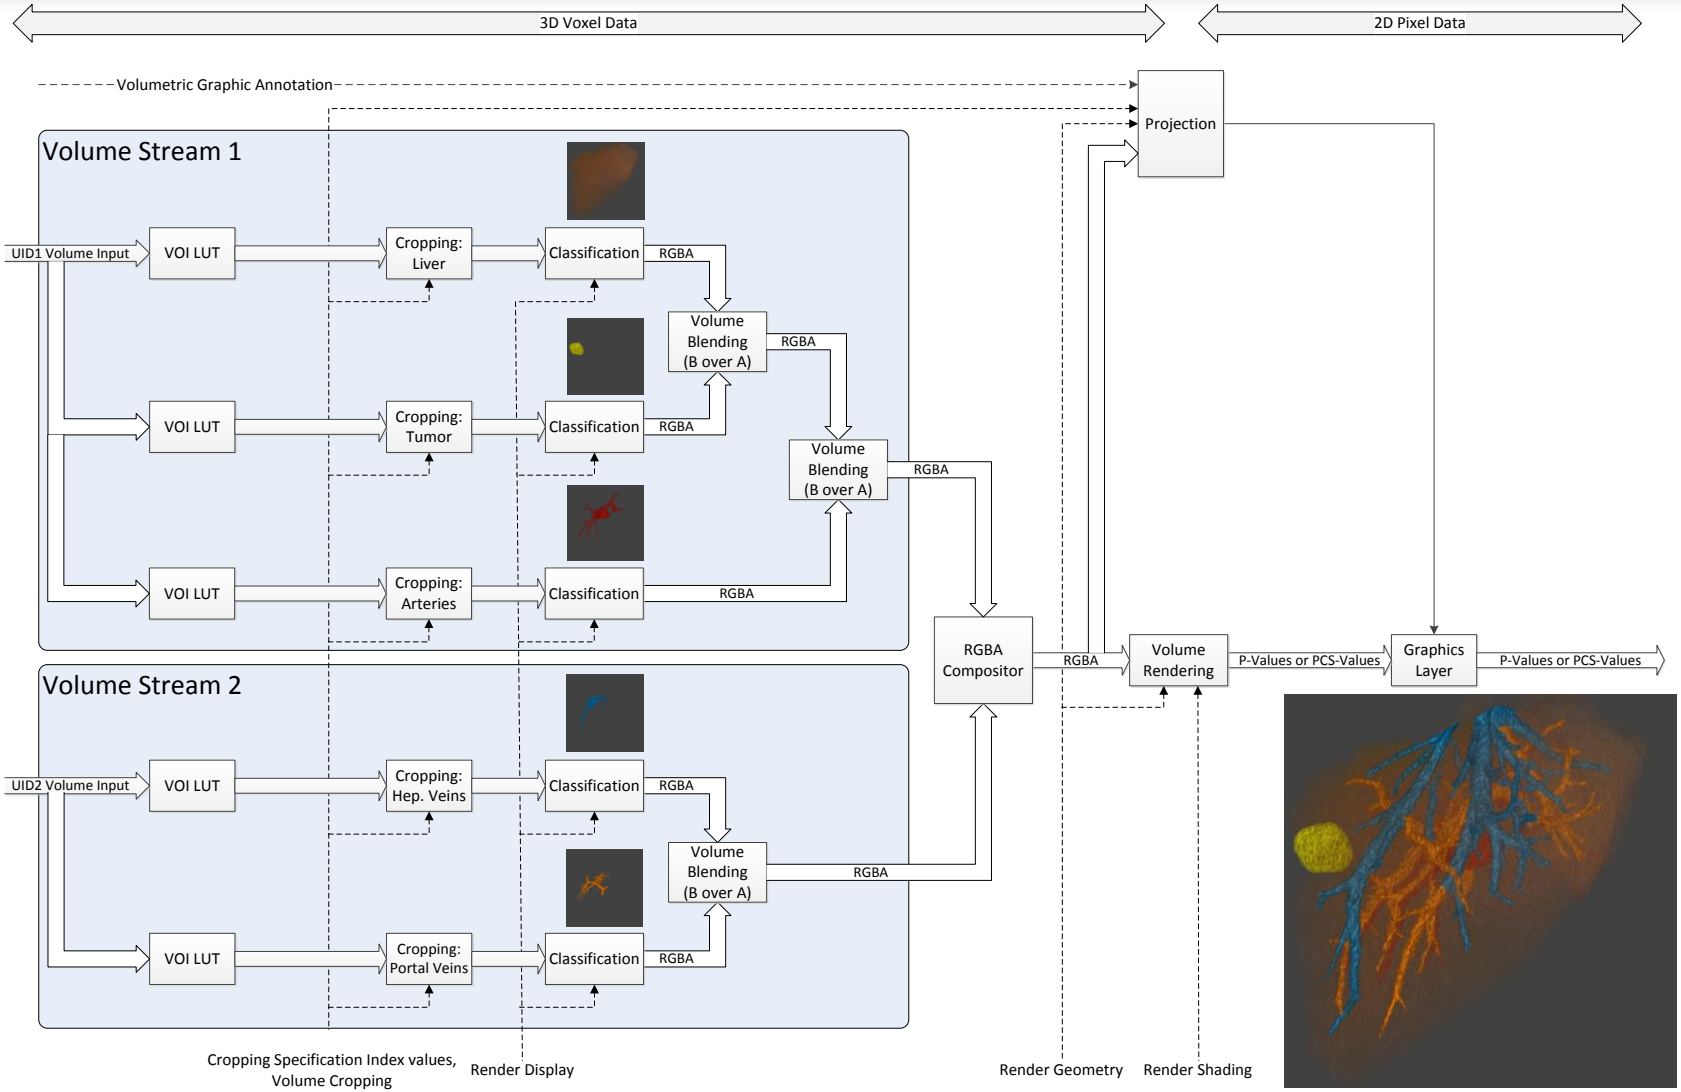

Leverages Planar MPR VPS

- Same basic concepts as Planar MPR VPS
  - Geometry based on a Reference Coordinate System
  - Modules from Planar MPR Information Object Definition:
    - VPS Identification to label the presentation
    - VPS Relationship to specify volume inputs
    - VPS Cropping to specify croppings that are applied to each input
    - Clinical Description to specify a clinical context of the view (anatomy, view name, etc.)
    - Graphic Annotation for placing graphics on 2D output view
    - Volume Graphic Annotation for placing graphics within the volume space to be rendered
    - Presentation Animation for animating the initial view

Adds New Layer to VPS Relationship Module

- Two-step Specification of inputs:
  - Volume Presentation Input Set Sequence
    - Specifies each input that can be shared among one or more items in Volume Presentation State Input Sequence
    - Each input set identified with Volume Presentation Input Set UID
  - Volume Presentation State Input Sequence
    - Instead of redundantly specifying the same input in multiple items, each item references an item in Volume Presentation Input Set Sequence
    - Uses Volume Presentation Input Set UID to specify the input set
    - This structure affects conformance of recently approved Planar MPR Volumetric Presentation States

Adds New Layer to VPS Relationship Module

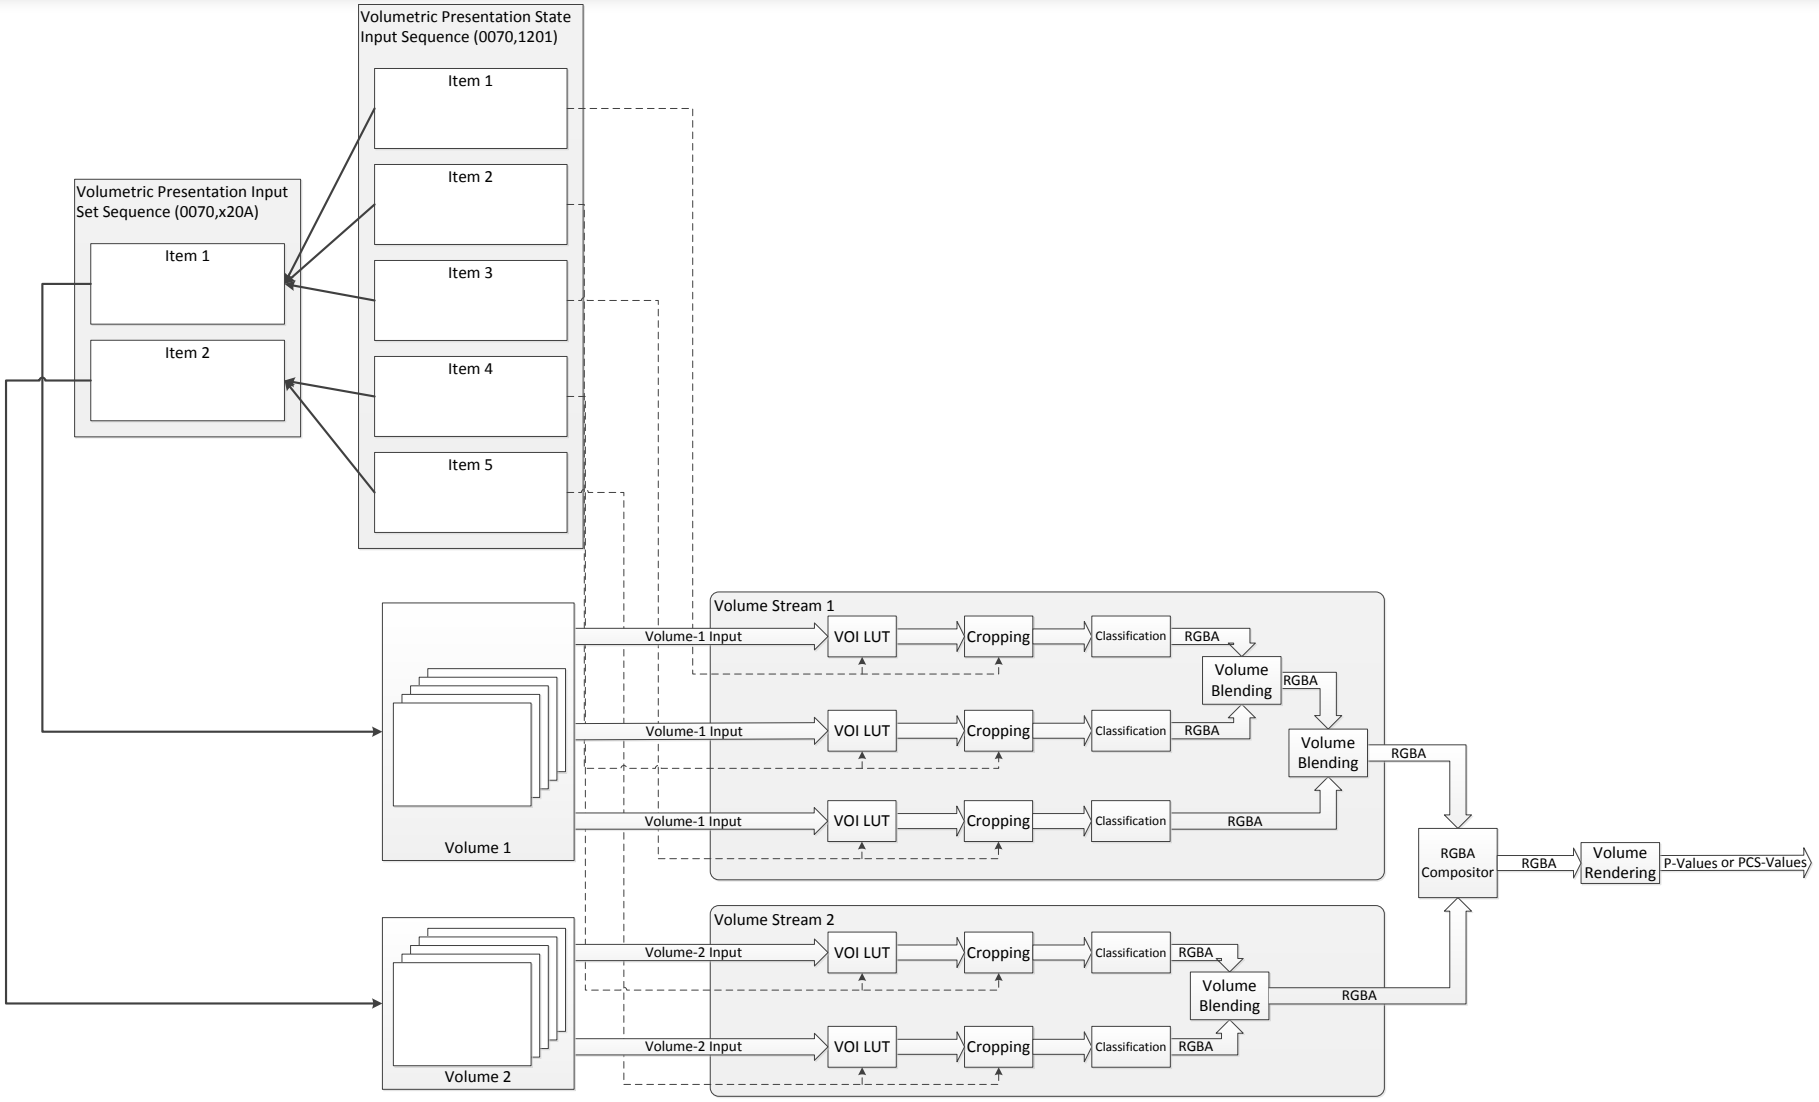

Volume Rendering

- Data visualization method used to display a 2D projection of a volume dataset
- Generally consists of the following steps:
  - Segmentation
  - Gradient Computation
  - Resampling
  - Classification
  - Shading
  - Compositing
- Significant latitude given to implementation-specific decisions on implementing each step

Segmentation

- Separating the volume data into groups that will share a particular color palette.
- **In-scope:** Segmentation objects are specified as cropping inputs to the Volumetric Presentation State.

Gradient Computation

- Finding edges or boundaries between different types of tissue in the volumetric data.
- Typically used to determine "surface normals" for use in the shading operation
- Gradient Computation used is an implementation decision outside the scope of the Volumetric Presentation State.

Resampling

- Resampling volumetric data creates new samples along an imaginary ray behind each pixel in the output two-dimensional view
- Generally uses some interpolation of the values of voxels in the neighborhood of each sample
- **In-Scope:** The geometry specifying the output view is specified in the Render Geometry module of the VPS
- The interpolation method used is an implementation decision outside the scope of the Volumetric Presentation State.

Classification

- Assigns a color and opacity to each ray sample
- **In-scope:** Classification parameters are specified in the Render Display module of the VPS

Shading

- Application of a lighting model to ray samples indicating the effect of ambient, diffuse, and specular light on each sample
- **In-scope:** Basic shading parameters are specified in the Render Shading module of the VPS.
- There is considerable research and development activity in this area as vendors strive to create the most "natural" looking image presentations

Compositing

- Accumulation of samples on each ray into the final value of the pixel corresponding to that ray
- The specific algorithms used are outside the scope of the Volumetric Presentation State.

Order of Volume Rendering Operations

- Reference pipeline implies a particular order of volume rendering operations
- Attributes are defined in terms of this reference pipeline
- An implementation is free to use any order of volume rendering operations
- The implementation is responsible for transforming the standard attributes into parameters appropriate for the particular ordering of operations used

New Modules for Volume Rendering VPS

- Render Geometry
- Render Shading
- Render Display

Render Geometry Module

- Specifies the render view (or initial render view if animation is used)
- Includes
  - Projection: Orthographic or Perspective
  - Viewpoint
    - Position
    - Look-at point
    - Up direction
  - Field of View (X_{left}, X_{right}, Y_{top}, Y_{bottom}, D_{near}, D_{far})
  - Compositing method (AveIP, MinIP, MaxIP, Render)

Orthographic Field of View Geometry

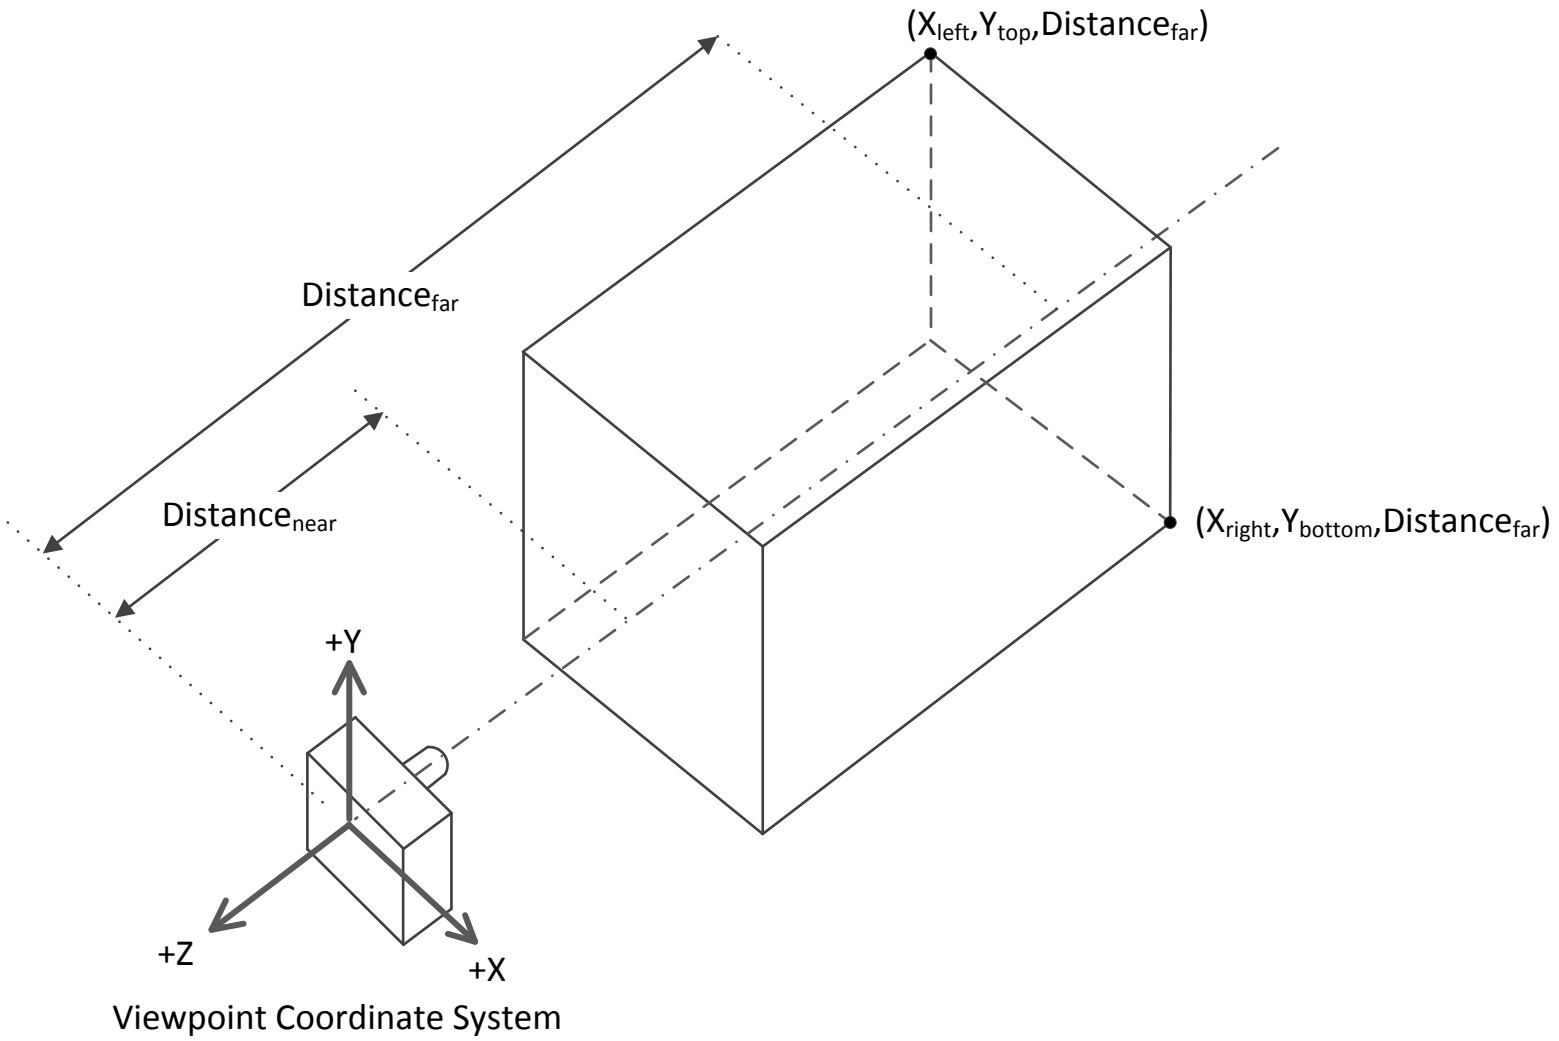

Viewpoint Coordinate System

Perspective Field of View Geometry

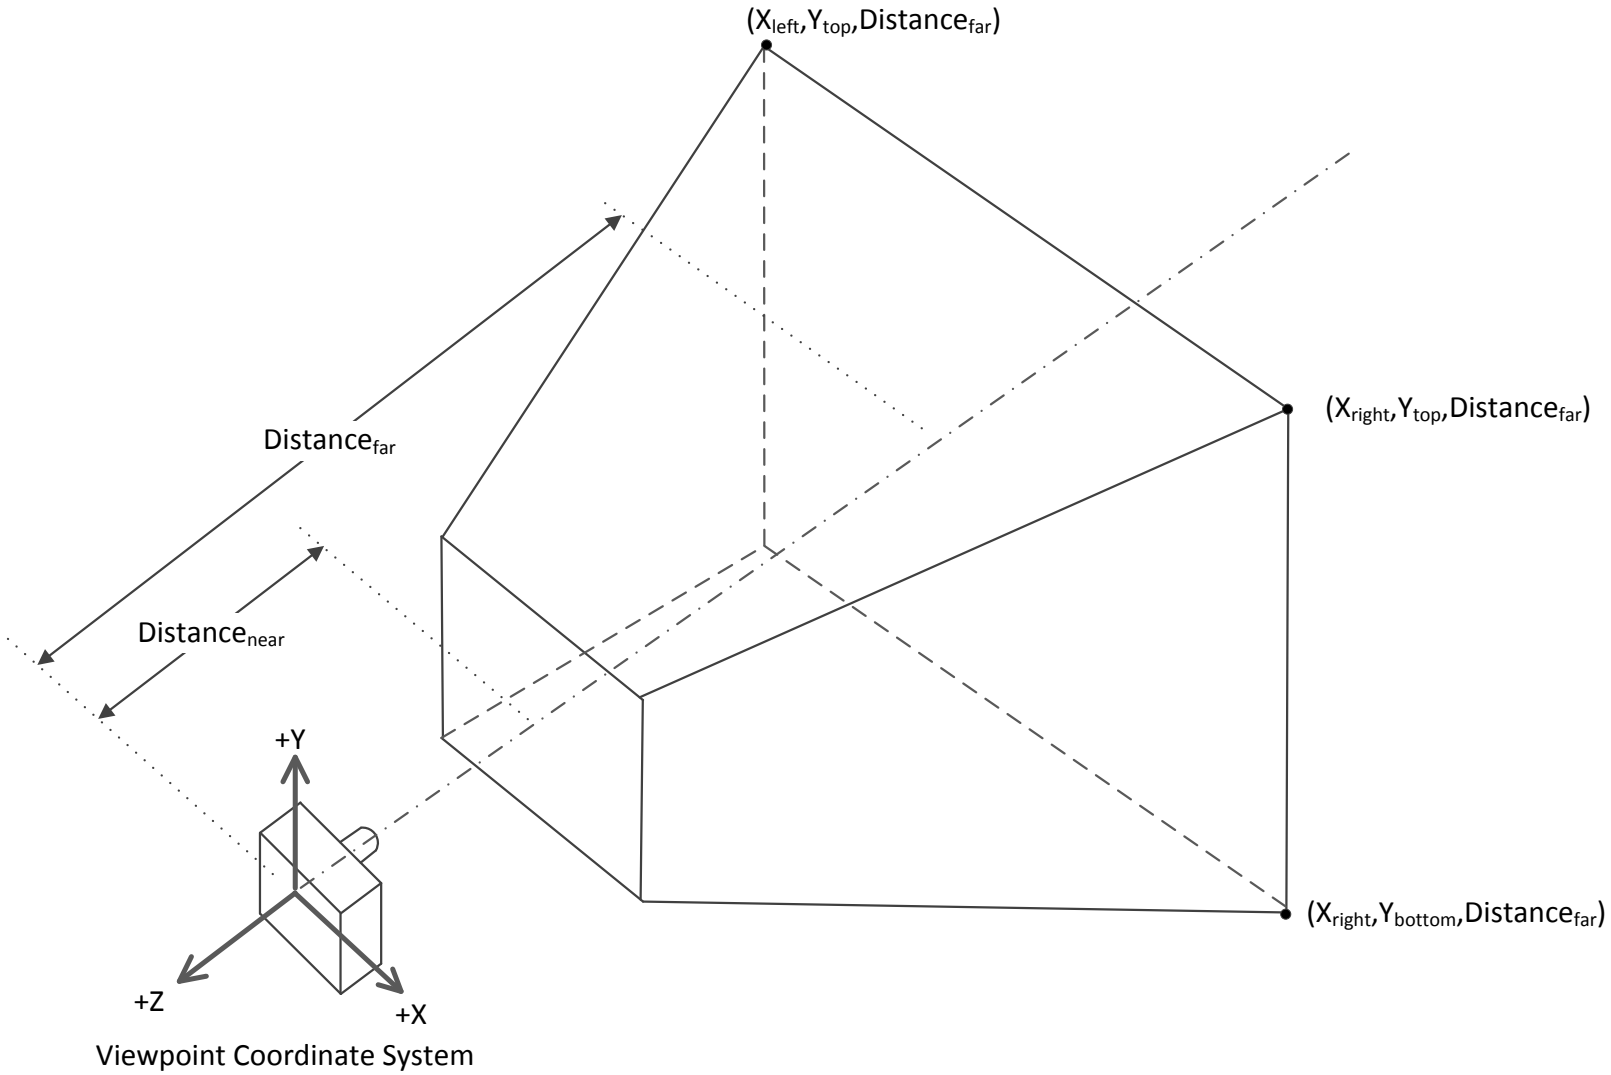

Render Shading Module

- Provides basic lighting parameters, assuming a single white light source at infinity for diffuse and specular lighting
- Includes
  - Shading Style (single-sided vs. double-sided)
  - Reflection Intensities
    - Ambient
    - Diffuse
    - Specular
  - Light Direction
  - Shininess
- Subject to extension by the rendering application

Render Display Module

- Each VPS input has separate RGB and Alpha (opacity) lookup tables (i.e., Classification)
- Inputs are grouped into "Volume Streams"
  - Segmented Volume Render VPS has one Volume Stream
  - Multiple Volume Render VPS has multiple Volume Streams
- Within each Volume Stream, inputs are blending using fixed "B over A" blending to produce a single RGBA output
- Volumes output from multiple Volume Streams are RGBA Composited prior to volume rendering

Alpha vs. Opacity

- Alpha and Opacity are treated as identical concepts
- The same specification of Alpha/Opacity is used for both
  - Compositing of multiple volumes prior to rendering (blending of volumes)
  - Compositing of samples on each ray contributing to a pixel in the output view (ray casting)
- No separate blending and ray casting opacity specifications

Rendering Pipeline

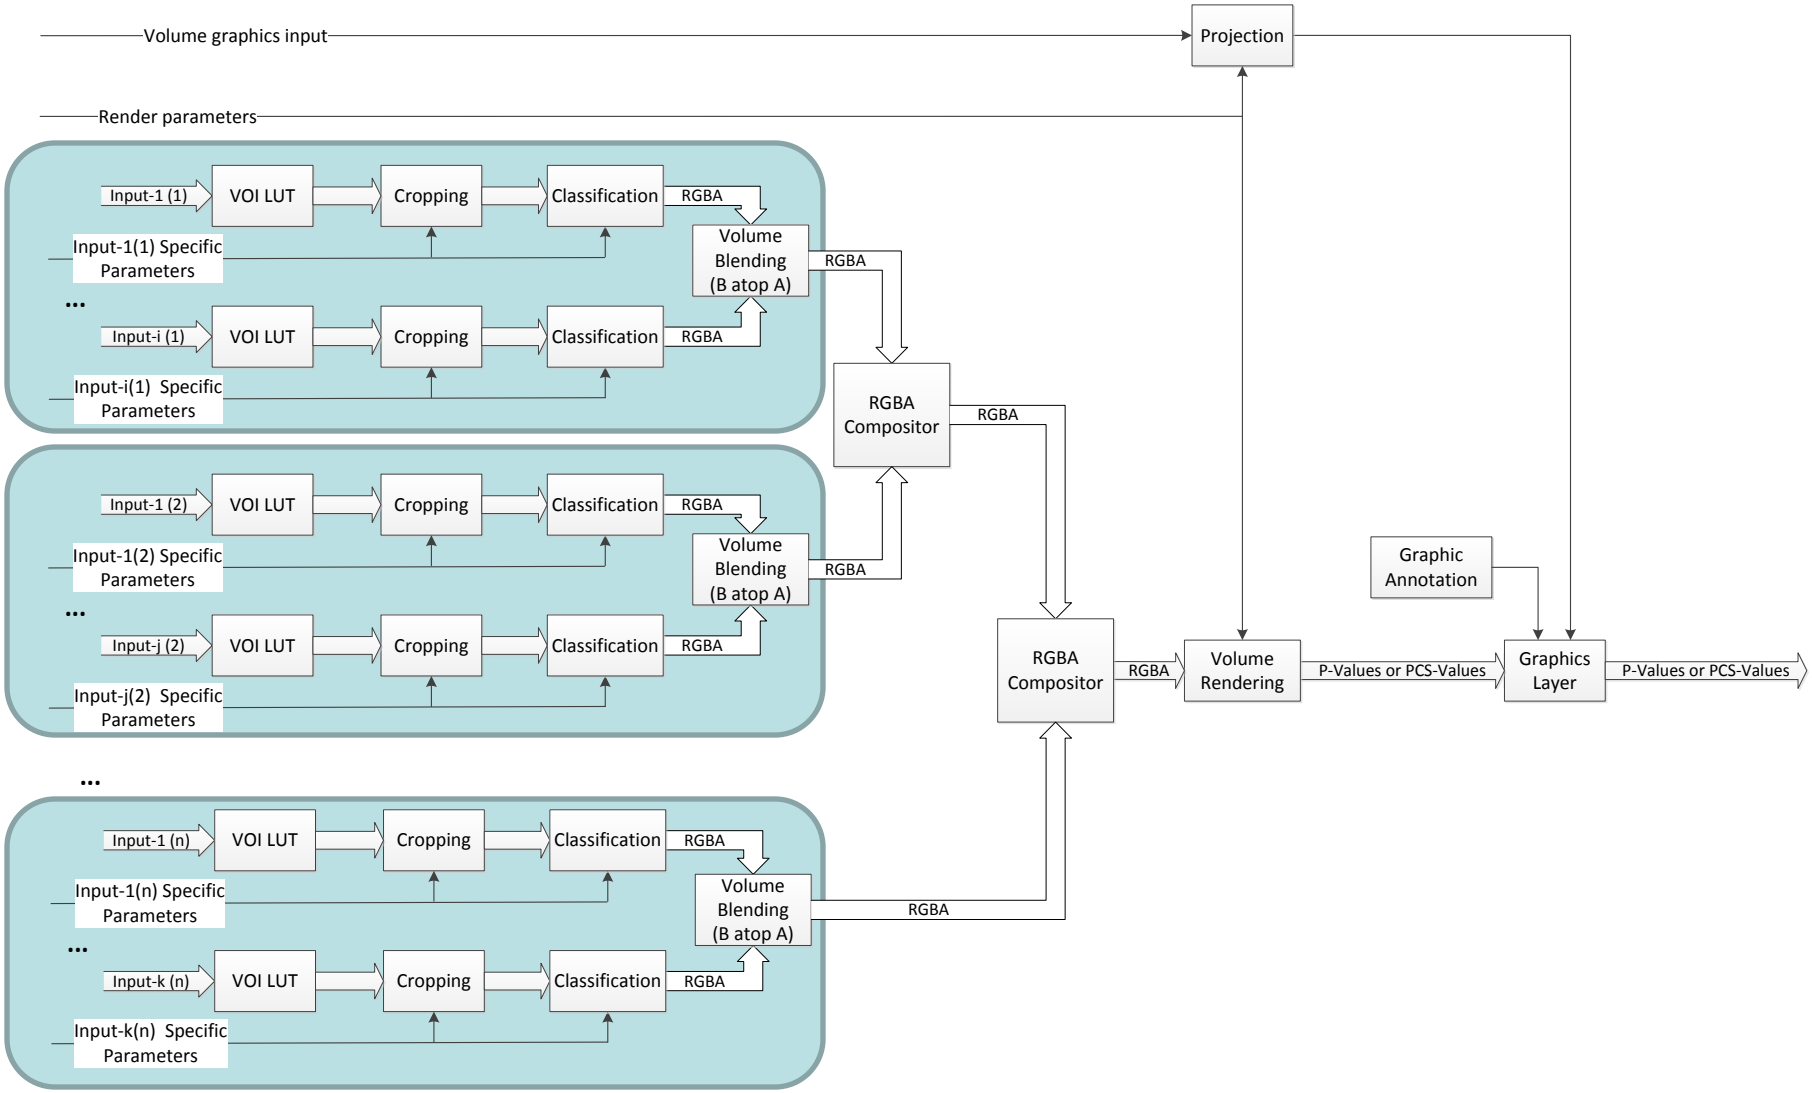

RGBA Compositing

- In the reference pipeline, the volume dataset is composited (i.e., converted to RGBA) prior to volume rendering
- Same flexible compositing structure used for Planar MPR, consisting of
  - Weighting Tables that consider Alpha inputs from both channels in determining the weighting factors for each RGBA input to the Compositor Function
  - Compositor Function that combines each RGBA component using the relationship:
    $C_{out} = (C_1 * Weight_1) + (C_2 * Weight_2)$

RGBA Compositor Structure

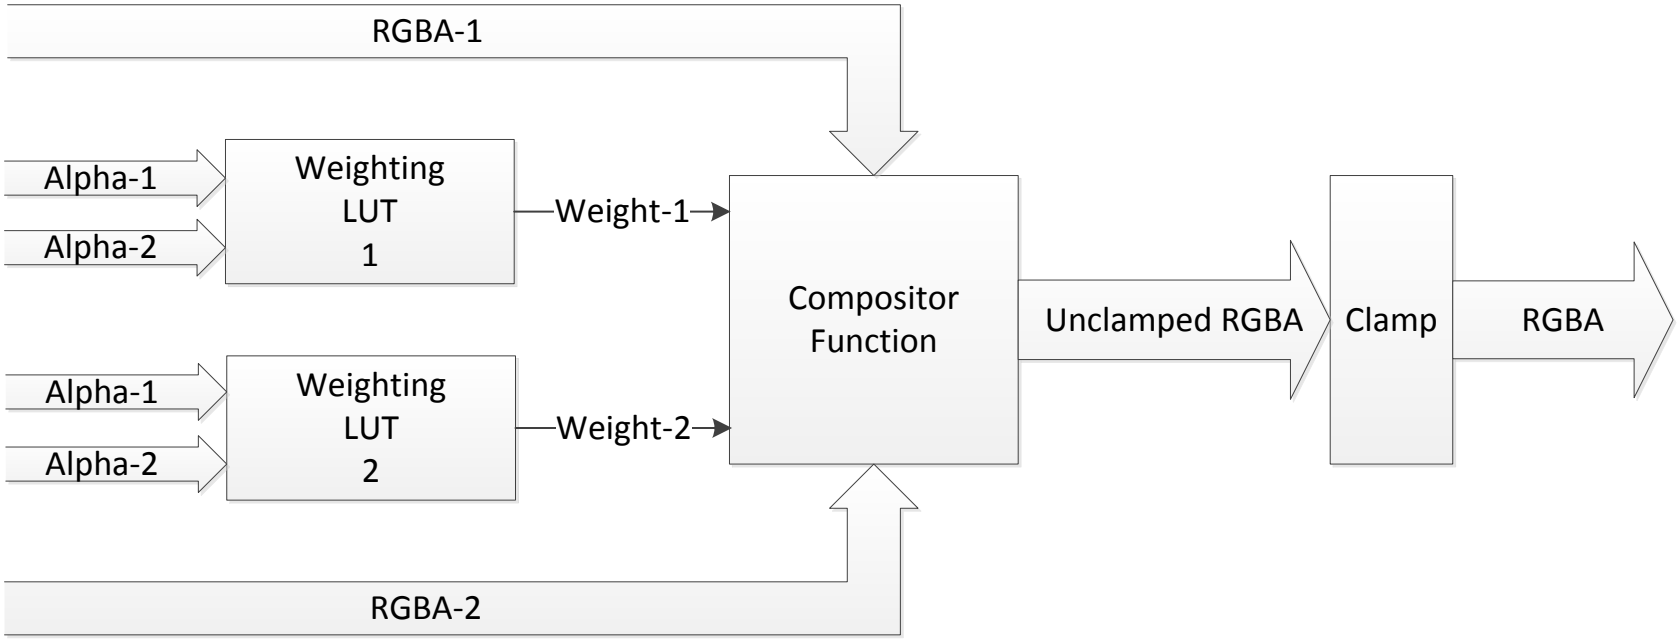

Volume Rendering Component

- Inputs to Volume Rendering component are
  - Single, blended, RGBA volume dataset
  - Render Geometry parameters
  - Render Shading parameters
- Certain rendering steps are left to the application, including
  - Gradient Computation
  - Resampling / interpolation
  - Compositing
  - More advanced (proprietary) algorithms

Volumetric Graphics Annotation Projection

- Inputs to Projection component are
  - Volume Graphic Annotation module
  - RGBA volume to be rendered
  - Input-specific Cropping Specification Index values
  - Volume Cropping module
- Projection transform creates graphics layers from the volumetric graphics annotation specification
- Annotation Clipping defines alternative behaviors:
  - If clipping, only show graphics if not obscured by opaque data in the foreground
  - If not clipping, show graphics always

Presentation Animation

- Two new animation styles are added
  - Flythrough: The viewpoint moves along a pre-defined curve and the view is re-rendered at each step. Effect is that of "flying through" the volume space while viewing the anatomy.
  - Swivel: The viewpoint remains fixed while the volume data rotates back and forth around a "swivel axis"

Flythrough Animation

- A set of discrete points defining a curve is provided
- Viewpoint look at point moves along the curve at a specified velocity (mm/sec)
- Viewpoint position is adjusted at each step to maintain the original distance to look at point
- Viewpoint up direction maintained according to the cross-product
  $V_{originalLook} \times V_{originalUp} = V_{look} \times V_{up}$

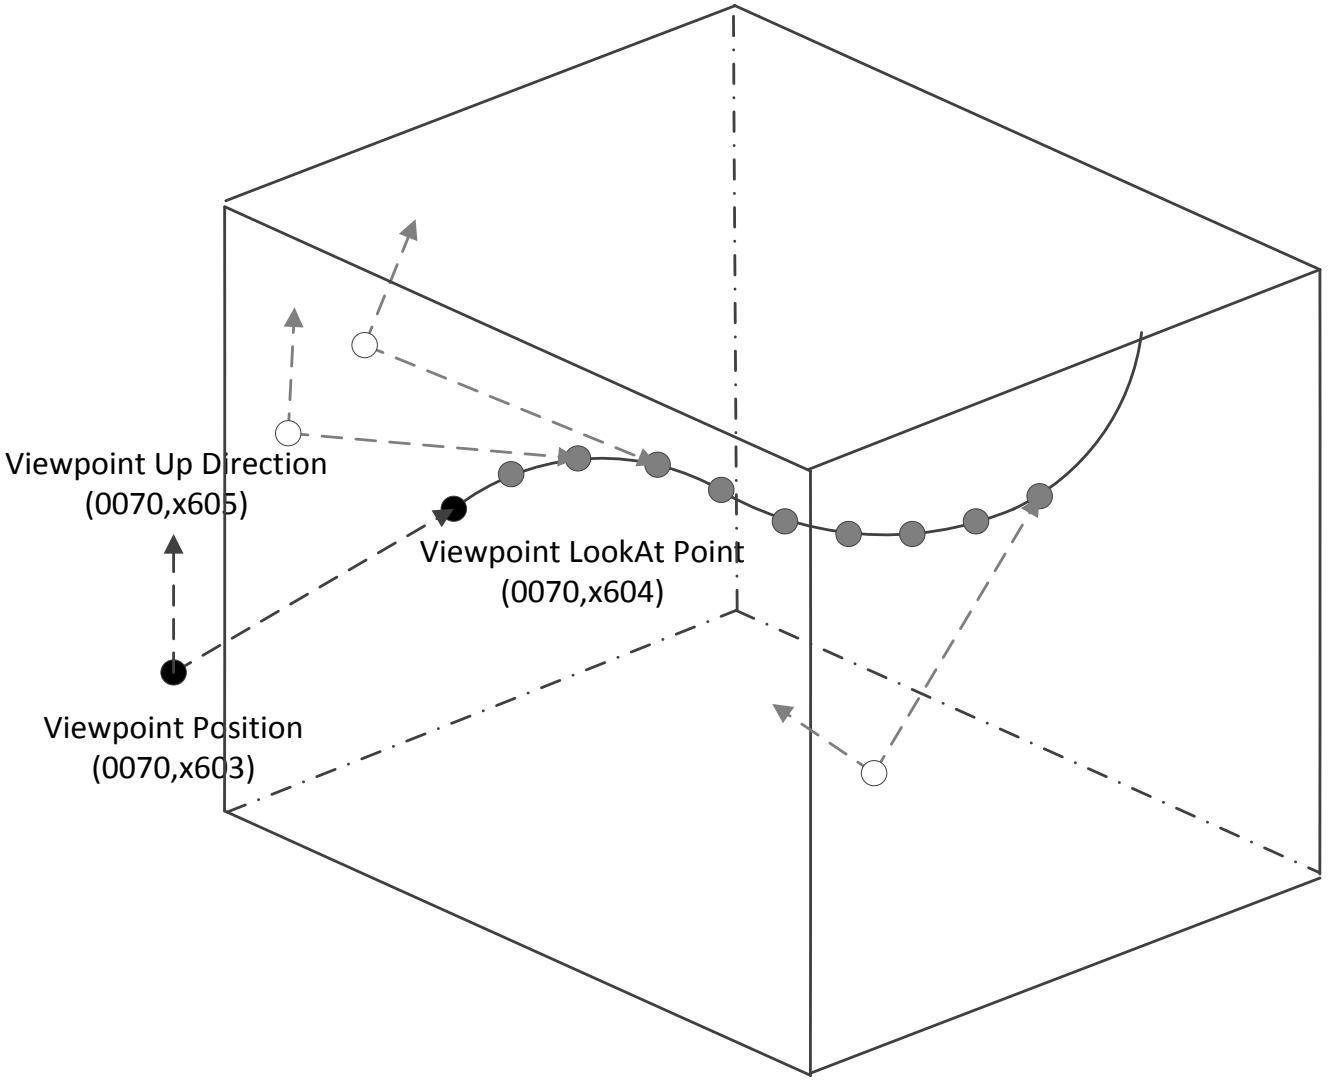
Viewpoint Up Direction
(0070,x605)
Viewpoint LookAt Point
(0070,x604)
Viewpoint Position
(0070,x603)

Swivel Animation

- Swivel axis defined as the axis parallel to Viewpoint Up Direction intersecting the Viewpoint Look At Point
- The rendered volume rotates back and forth in the specified angular range at a specified rotational frequency
- Currently, no guidance is provided on a need for smooth direction changes (e.g., sinosoid rotational velocity vs. constant)

Output

- Final View may be either
  - Grayscale (P-Values)
  - Color (PCS-Values)
- As with Planar MPR, may include links to a "secondary capture" of the presentation at the source device
  - For quality assurance
  - For clinical integrity