

Volume Rendering Assignment

2IMV20

Technische Universiteit Eindhoven

Patrick Strookappe 0896169

Jakob Serlier 1645757

Wouter van Oorschot 0906557

Group 67

December 13, 2020

1 Introduction

Visualisation methods are an increasingly important topic of research due to their importance in fields like medical imaging. In this report we will discuss multiple classical and state of the art visualisation techniques. We start by discussing the theoretical foundation of each technique. Then we discuss our implementation and show how our results compare to the examples in the assignment and demo video. Finally in the Analysis we discuss the advantages of each technique and compare them with each other using interesting visualisations.

2 RayCasting

2.1 Tri-linear interpolation

Instead of the standard nearest neighbour approach to retrieve voxel values, we use tri-linear interpolation for all visualisation techniques.

Similarly to `getVoxel()`, the first step of the interpolation checks whether or not the coordinates are inside the volume. However, because we now also use `ceil()` values of the coordinates, an index out of bounds exception might occur. This coordinate-inside-volume check was slightly adjusted (by `[-1,-1,-1]`), which prevents index out of bounds at the extremes.

Next, the 8 vertices of integer coordinates, representing a surrounding cube, are calculated by taking the `floor` and `ceil` values of the original 3 coordinate points. These points of the cubic lattice are labeled v_{000} to v_{111} .

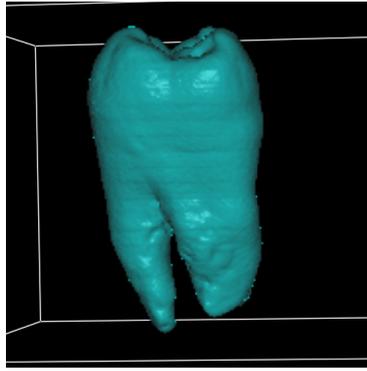
Using the tri-linear formula described in the lectures we can now calculate the desired interpolated value:

$$\begin{aligned} val = & (1 - \alpha)(1 - \beta)(1 - \gamma)V_{000} \\ & + (\alpha)(1 - \beta)(1 - \gamma)V_{100} \\ & + (1 - \alpha)(\beta)(1 - \gamma)V_{010} \\ & + (\alpha)(\beta)(1 - \gamma)V_{001} \\ & + (1 - \alpha)(1 - \beta)(\gamma)V_{110} \\ & + (\alpha)(1 - \beta)(\gamma)V_{101} \\ & + (1 - \alpha)(\beta)(\gamma)V_{011} \\ & + (1 - \alpha)(1 - \beta)(\gamma)V_{111}. \end{aligned} \tag{1}$$

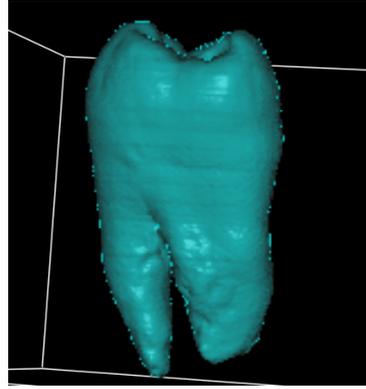
For the formula above, α , β , and γ represent the distance between the coordinate values and the rounded floor values. Figure 1a shows a subtly sharper image, with less artifacts, as opposed to 1b, confirming that our implementation is correct. Further methods and images shown below use tri-linear interpolation where possible. Tri-linear interpolation for gradients is discussed in section 3.2.1.

2.2 Compositing ray functions

The compositing ray function is created using the same basis as the MIP-raycasting function. For each step on the rayvector the voxel value is calculated using tri-linear interpolation. The transfer function then maps this value to a final value, based on the parameters we set in the transfer function. All these transfer function colors along the ray are added using the front-to-back version of the compositing function from the 1995 optical model paper [5],



(a) Tooth dataset using tri-linear interpolation, isosurface raycasting (ISO 500) and Phong shading: little artifacts.



(b) Tooth dataset using nearest neighbour, isosurface raycasting (ISO 500) and Phong shading: increased artifacts.

Figure 1: Difference between trilinear interpolation and nearest neighbour, with isosurface raycasting and phong shading.

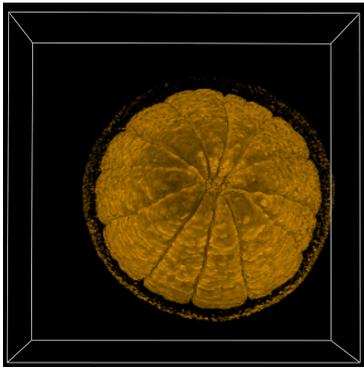
which is simplified in the lecture to equation 2. In the implementation (equation 3) we calculate each new color C_i in the ray using a color c_i , based on the voxel color and transfer function, the current opacity and the previous color C_{i-1} . Additionally, α_i is the opacity of the current voxel and calculated using equation 4, which we get a from the transfer function. Equation 4 uses a similar approach to equation 3 where it iteratively calculated the new α based on the previous value.

$$I(p) = \sum_{i=0}^{n-1} c_i \prod_{j=i+1}^{n-1} (1 - \alpha_j). \quad (2)$$

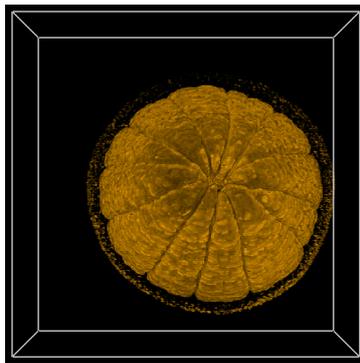
$$C_i = \alpha_i * c_i + (1 - \alpha_i) * C_{i-1} \quad (3)$$

$$\alpha_i = \alpha_{i-1} + a_i * (1 - \alpha_i) \quad (4)$$

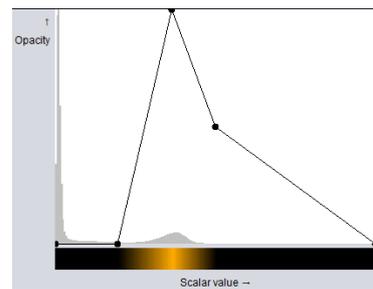
The composite function is compared to the examples given in the assignment. In this case the Orange8 dataset. As shown in image 2 the results are visually nearly identical.



(a) Drawn composite image of the orange



(b) Composite image from the assignment



(c) Transfer function for figure 1b

Figure 2: Orange comparison

2.3 Resolution optimizations

The raycasting calculations are too heavy to allow for a responsive user experience when moving or rotating the object by default. Because of this a global boolean `InteractiveMode` indicates if the user is manipulating the view area. If this is the case, both the `sampleStep` and the `increment` value are changed from 1 to 3. The `sampleStep` value lowers the amount of steps calculated over a casted ray. The `increment` value determines the amount of rays casted per pixel on screen. When this is increased, the output resolution of the image is lower. This results in a fluent user experience, but still clear enough to be able to see how the object is rotated or translated.

3 Isosurface Raycasting and Shading

3.1 Isosurface Raycasting

The implementation of `traceRayIso` largely follows the existing implementation of `traceRayMIP`. Following the ray which is cast, we determine the increments and number of samples and initialize a current position. This current position changes when we move along the raycast. Next, as the current position starts moving along the ray in steps, the voxel-value is compared to the `isoValue` (`isoValueFront` or `isoValueBack`). When the voxel-value is greater than the user-determined `isoValue`, the voxel needs to be drawn and the opacity is set to 1, and iterating over the position stops. Then the isocolor is pulled from the UI. Allowing the isovalue to change is therefore what allows for interactive rendering of different sections of the volume.

The above method follows isoSurface rendering using raycasting, in contrast to Marching Cubes described in the lectures. Figure 3 shows a comparison between our results and the result from the assignment, confirming the correctness of isosurface raycasting, in addition, figure 7 shows the correctness of isosurface raycasting with Phong Shading enabled.

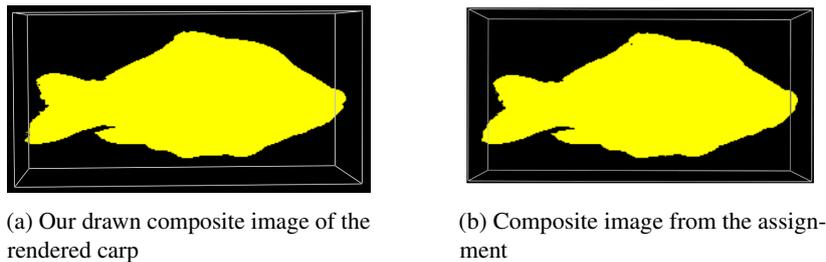


Figure 3: Iso surface raycasting carp comparison

3.2 Computing gradient

We use the gradient approximation from Levoy [2] as seen in equation 5. We do not calculate these values for the outermost edge voxels, as that prevents out of bound exceptions. These voxels are set to 0. This generally does not affect the visualisation as these edges generally are empty anyway. This formula was implemented in the `Compute()` function in `GradientVolume.java`.

$$\begin{aligned} \nabla f(\mathbf{x}_1) = \nabla f(x_i, y_j, z_k) \approx & \\ \frac{1}{2} \{ f(x_{i+1}, y_j, z_k) - f(x_{i-1}, y_j, z_k) \} & \\ \frac{1}{2} \{ f(x_i, y_{j+1}, z_k) - f(x_i, y_{j-1}, z_k) \} & \\ \frac{1}{2} \{ f(x_i, y_j, z_{k+1}) - f(x_i, y_j, z_{k-1}) \} & \end{aligned} \quad (5)$$

3.2.1 Tri-linear application to gradient

For tri-linear interpolation of gradients the approach is very similar to that of voxels. In tri-linear interpolation, a single final value is interpolated from 8 surrounding voxels. With gradients each of these voxels is a vector of size three. The only difference in these values is that they can be negative, which is not a problem in the formula. We apply the formula three times, separately for the gradient value of the coordinates of x, y, and z. This new coordinate value is returned as a `VoxelGradient` vector. This implementation is done in `getGradientTrilinear`.

3.3 Phong shading

Shading is the use of light reflection to improve the look of images. More than that, it gives both a sense of depth, and, when used on different materials, gives us clearer boundaries and helps differentiate between them. We implement a simplified version of Phong shading based on using a halfway vector as an approximation of the reflection vector. Examples of the implementation can be seen in figures 4, 5, and 7. The halfway vector (equation 6) avoids the expensive vector multiplications and dot products of the reflection vector (equation 7) by using simple vector additions. The usage of the `phongShading` is simplified further by assuming that the lightvector (L) originates in the same point as the Viewvector (V), such that $L = V$. The `phongShading` function does allow for a different light- and viewvector, but the visualization will always call it with $L = V$.

$$H = \frac{L + V}{|L + V|} \quad (6)$$

$$R = (2N \cdot L)N - L \quad (7)$$

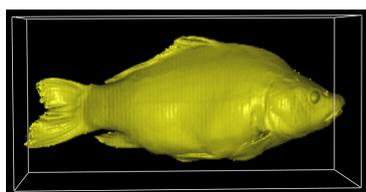
The complete implementation of the implementation is seen in equation 8. This calculation is done separately for the three colors (I_d).

$$I = I_a * k_{ambient} + I_d k_{diff}(L \cdot N) + k_{spec}(N \cdot H)^\alpha \quad (8)$$

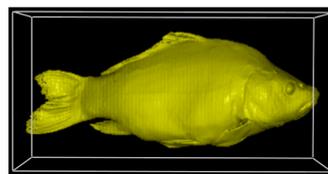
- I_a = Opacity of voxel.
- I_d = R, G, B value of voxel.
- k_{diff} = spectral diffusion
- $L = V =$ Viewvector
- N = normal vector = $\frac{\nabla(f(x_1))}{|\nabla(f(x_1))|}$
- k_{spec} = spectral diffusion
- $\alpha = n$ = exponent used to approximate highlight.

Which is a small alteration to the simplified model in the lecture slides. We can directly implement shading in Isosurface Raycasting by using $I_a = 1$. For Composite ray functions it is a bit more difficult as we need to shade each step in the ray and combine these values into one final RGB value. We use an adapted version of equation 3, where instead of color we keep track of a phongcolor and instead of a voxelcolor we calculate a phongcolor for every step, this is written down in equation 9. The α_{pix} value is the opacity value of the color of the previous step. For 2D transfer function we also need to update the Phong value every step. For this equation 9 is also used. In each step the the Phong color of the previous step (PhongColorPrevious) is used. In 2DTransferFunction the α_{pix} value is the alpha value from the previous step, times the opacity based weighting, as explained in the next chapter. [4].

$$PhongColor = PhongColor * \alpha_{pix} + PhongColorPrevious * (1 - \alpha_{pix}) \quad (9)$$



(a) Our drawn composite image of the rendered carp with Phong Shading



(b) Phong shading image from the assignment

Figure 4: Phong shading with Isosurface raycasting

4 2D transfer function: Gradient-based opacity weighting

Gradient based opacity weighing is an example of a separable 2D Transfer function. It uses the voxel value and it's gradient magnitude to create 2 1D functions, which can be combined to create a picture in which it is easier to classify different materials or features [3]. Gradient magnitude gives an idea of how quick changes are happening between different voxels. Such regions with rapid changes are useful to visualise as they may indicate a change in materials [1]. Our approach to combining the voxel values and gradients is presented in [2]. We use the same ray vector as in MIP and compositing ray functions. For each sample in this ray we want to compute an opacity. The intuition behind deciding on an opacity value is that we want to visualise a voxel if it's value $f(x_i)$ is the same as a value we have chosen: f_v in the formula and intensity in the transfer function. In this case we make it fully opaque. When the intensity is close to the voxel value we also want the opacity to be close, that's why we measure the distance between them $|f_v - f(x_1)|$. When gradients change edges become visible. We smoothen these by dividing by the gradient magnitude $|\nabla f(x_1)|$. Finally we also only calculate these opacities for values that are within some radius r , such that we can decide which voxels we want to display. In the implementation this r is scaled by the maximum gradient. If voxels are not in this radius they are set to completely transparent. Similarly to composite ray functions, we calculate the final opacity by iteratively combining the opacity of the previous voxel in the rayvector and the newly retrieved opacity from the transfer function. This approach is described in Levoy [2] (equation 11). In the Implementation this value α_{pix} is calculated iteratively each step of the rayvector.

$$\alpha(x_1) = \alpha_v \begin{cases} 1 & \text{for } |\nabla f(x_1)| = 0 \text{ and } f(x_1) = f_v \\ 1 - \frac{1}{r} \frac{|f_v - f(x_1)|}{|\nabla f(x_1)|} & \text{for } |\nabla f(x_1)| > 0 \text{ and } f(x_1) - r|\nabla f(x_1)| \leq f_v \leq f(x_1) + r|\nabla f(x_1)| \\ 0 & \text{otherwise} \end{cases} \quad (10)$$

$$\alpha_{\text{pix}}(x_1) = 1 - \prod_{n=1}^N (1 - \alpha_n(x_1)) \quad (11)$$

5 Cutting plane

The first challenge with the cutting plane assignment is determining the side of which a voxel is in regards to the cutting plane. In order to do so, we created a helper function `leftOfPlane` which uses an altered version of the point-normal formula of a plane in order to return a Boolean that determines whether the voxel is left of the plane or not. In this function, the `planeNorm` and `planePoint` are used to determine the side of the voxel:

$$a(x - x_c) + b(y - y_c) + c(z - z_c) + d > 0. \quad (12)$$

The values a , b , c and x , y , z , are taken from the already implemented variables `planeNorm` and `planePoint` respectively. x_c , y_c and z_c are the coordinates of the point-coordinate to be checked. This allows us to check if a point is left or right of a plane, representing the front and back of the image.

Next, in the raycast function, the cutting plane logic is implemented by calling the `leftOfPlane` helper function on the current pixel coordinate. Instead of a single switch statement to calculate the color values, two are now used, for the front and back modes respectively. A helper variable `isFront` is created and used which is used in the various modes to determine if we should use the values from the front or back modes (for instance, `isoValueFront` or `isoValueBack`). This helper variable is then passed as an argument in the method functions.

Specific mode implementations:

In the `traceRayIso` function, the helper variable `isFront` is used to use either `isoValueFront` or `isoValueBack`, and to use the `rgb` values of the `isoColorFront` and `isoColorBack`. Since it already returns a value which is either positioned left or right of the plane, this concludes the implementation of cutting plane for `traceRayIso`. The modes `COMPOSITING` and `TRANSFER2D` are a bit more complex. They are contained in the `traceRayComposite` function. Here, `isFront` is used to take both the correct voxel values and material values (where required) depending if the voxels are left or right of the plane.

The switch statement for the modes within `traceRayComposite` needed to be updated to contain the truth value of the specific `isFront` value for the voxels that are being calculated. This was done by creating a new `Raycast Mode` instance which changes depending on the `isFront` variable. Therefore, both the modes in the switch statements can now correctly use the color and material values for the front and back modes respectively, which completes the implementation for cutting plane of each of the rendering modes.

6 Analysis

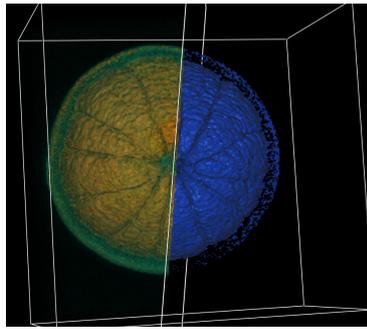
Comparison of modes: The modes `MIP` and `slicing` were pre-implemented, and will therefore not be a major part of this analysis. The current implementation of `slicer` is useful when taking a look at a single slice of an object. `MIP` allows for a quick identification of high density objects within objects.

Identifying use-cases and comparison of composite and 2d transfer function modes: In image 7a we can see a comparison between the 2d transfer function and the 1d transfer function (compositing mode). When attempting to find the coins hidden inside the pig dataset, from a practical perspective, the 2d transfer function allowed us to pinpoint the coins a lot easier and quicker. However, the 1d transfer function allows for the use of more colors, making us able to better distinguish the coins. When attempting to distinguish multiple objects within the dataset with varying densities ('values'), the composite function is especially useful by giving them different colors to better distinguish them as seen in figure 7b.

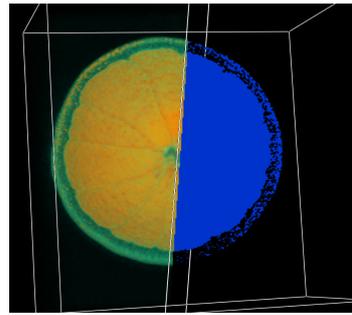
Comparing compositing raycasting to the `MIP` function, we see in figure 8 it gives more detail and more control over what can be seen. Using a transfer function, different intensity values of voxels on the ray can be given different colors and intensities. Using this, internal objects (such as organs) within other objects can become more clearly visible, whereas `MIP` only returns the most intense (highest voxel value) value throughout the ray.

Figure 5 clearly illustrates the added value Phong shading gives. It allows us to see depth in an image giving new information which was not seen before. Another image which clearly illustrates the advantages gained is the example from the isosurface mode on the pig: images 6a and 6b. Without Phong shading, the flower pattern could not be distinguished.

Finally, figure 9 showcases how the cutting mode allows two distinctly different modes to work well together. In the tooth example, the isosurface mode allows the viewer to get a clear understanding of the depth and shape of the tooth, where the 2d transfer function gives a solid impression of the internals of the tooth. All in all, different modes allow for different methods of analysis. Each of the modes can find added value for different datasets, and combining techniques (especially using the cutting plane) gives for the most broad analysis.



(a) Orange dataset: Phong enabled, composite and 2d transfer function

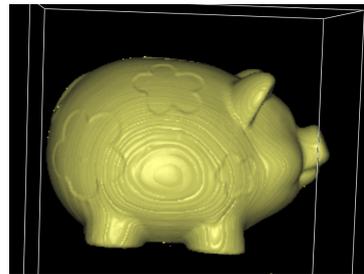


(b) Orange dataset: Phong disabled, composite and 2d transfer function

Figure 5: Showcasing composite and 2d transfer modes with Phong shading enabled and disabled

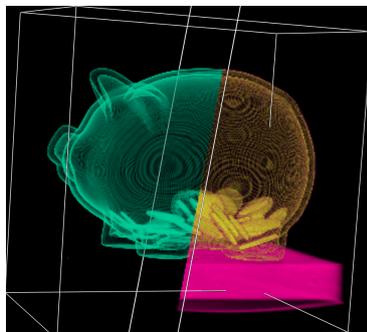


(a) Pig dataset with default isosurface raycasting settings

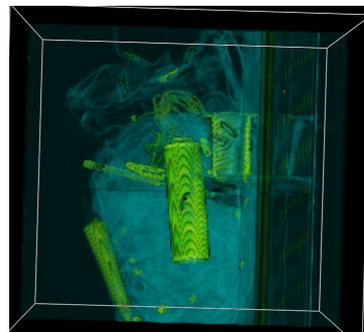


(b) Pig dataset with default isosurface raycasting settings and phong shading enabled

Figure 6: Finding hidden features using phong shading

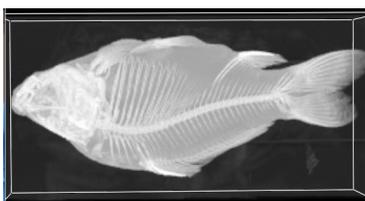


(a) Pig dataset. Left slice: 2d transfer function. Right slice: compositing mode

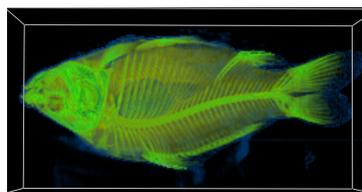


(b) Backpack dataset: use-case for multiple colors in composite function

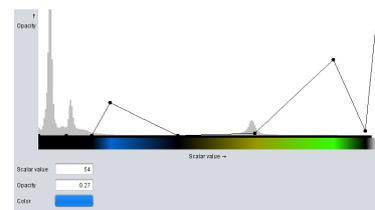
Figure 7: Practical difference in analysis 2d transfer function and compositing (1d transfer function) modes allow



(a) Carp dataset with MIP

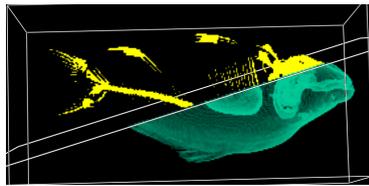


(b) Carp dataset with composite mode

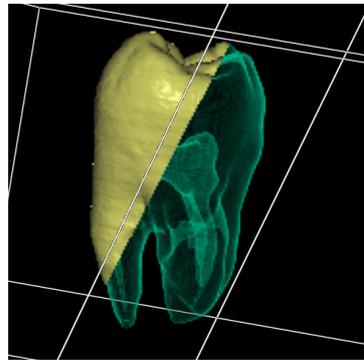


(c) Transfer function for figure 8b

Figure 8: MIP versus composite mode



(a) Carp dataset with isosurface raycasting (ISO 250) and 2d transfer function



(b) Tooth dataset. Left: isosurface (ISO 500). Right: 2d transfer function (DEFAULT). Phong shading enabled

Figure 9: Isosurface raycasting versus 2d transfer function: depth and shape versus internals

References

- [1] J. Kniss, G. Kindlmann, and C. Hansen. Multidimensional transfer functions for interactive volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 8(3):270–285, 2002.
- [2] M. Levoy. Display of surfaces from volume data. *IEEE Computer Graphics and Applications*, 8(3):29–37, 1988.
- [3] Patric Ljung, Jens Krüger, Eduard Gröller, Markus Hadwiger, Charles Hansen, and Anders Ynnerman. State of the art in transfer functions for direct volume rendering. *Computer Graphics Forum*, 35:669–691, 06 2016.
- [4] E.B. Lum and kun ling Ma. Lighting transfer functions using gradient aligned sampling. volume 2004, pages 289– 296, 11 2004.
- [5] N. Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995.

