

Computer Graphics


Dr. rer. nat. Martin Möddel

April 6, 2021

Institut für Biomedizinische Bildgebung

Ray Tracing

Ray Tracing



Ray Tracing



Definition

A ray tracer is an IOR algorithm, where a ray is traced into a scene until it hits a scene object. Information on hit point and object are used to color image pixels. A basic ray tracer has three parts

1. Ray generation: compute origin and direction of pixels viewing ray based on camera geometry and orientation
2. Ray intersection: Find closest object intersecting the viewing ray
3. Shading: Compute pixel color based on the result of the ray intersection

Definition

A ray is a 3D parametric line

$$p(t) = o + t(s - o). \quad (1)$$

Here o is the origin of the ray, $d = s - o$ is the direction and $t \in \mathbb{R}$ is the fractional distance of $p(t)$ to o .

- $p(0) = o$ and $p(1) = s$
- If $0 < t_1 < t_2$ then $p(t_1)$ is closer to the origin o than $p(t_2)$
- If $t < 0$ then $p(t)$ is behind o

Remark

Ray generation is tied to the camera and is usually performed in the camera coordinate system.

Ray Intersection

Apart from ray generation, intersection methods are an integral part of ray tracing, though they find their use in many other applications in computer graphics. Intersection methods should answer the following questions:

- Is my ray hitting any object?
- Are the objects in front of my camera, on its back or surrounding the camera?
- Which object is closest in front of the camera?
- Where does the object get hit?
- Where does the normal of the surface points towards?

Ray Sphere Intersection

Using the implicit formula we can describe a sphere centered at $c \in \mathbb{R}^3$ with radius r by $f(x) = \|x - c\|_2 - r = 0$. Let $p(t) = o + td$ be a ray then

$$\begin{aligned}0 &= \|o + td - c\| - r \\ \Leftrightarrow r &= \|o + td - c\| \\ \Leftrightarrow r^2 &= \|o + td - c\|^2 \\ \Leftrightarrow 0 &= t^2 + t \frac{2d \cdot (o - c)}{d \cdot d} + \frac{(o - c) \cdot (o - c) - r^2}{d \cdot d} \\ \Leftrightarrow 0 &= t^2 + 2tb + c,\end{aligned}$$

where $b = \frac{d \cdot (o - c)}{d \cdot d}$ and $c = \frac{(o - c) \cdot (o - c) - r^2}{d \cdot d}$.


Ray Sphere Intersection

The possibly complex roots are then given by

$$t_{\pm} = -b \pm \sqrt{b^2 - c}. \quad (2)$$

- If $b^2 - c < 0$ then the roots are complex. Meaning the ray does not intersect with the sphere.
- If $b^2 = c$ then the two roots coincide and the ray intersects with the sphere at one point.
- If $b^2 - c > 0$ then there are two intersection points.
 - If $b > \sqrt{b^2 - c}$ then the sphere is behind the camera
 - If $b < -\sqrt{b^2 - c}$ then the sphere is in front of the camera
 - If $\sqrt{b^2 - c} > b > -\sqrt{b^2 - c}$ then camera is inside the sphere

Ray Sphere Intersection - Optimizations



- Compute the vector $l = c - o$ and its length $l^2 = l \cdot l$
- If $l^2 < r^2$ the ray origin is inside the sphere (right)
- Compute $s = l \cdot d$
- If $s < 0$ and the origin is outside the sphere the sphere is behind the camera
- Compute distance between ray and sphere center $m^2 = l^2 - s^2$
- If $m^2 > r^2$ the sphere is going to be missed
- If not intersections occur at $t_{\pm} = s \pm q$, where $q^2 = r^2 - m^2$

Ray OBB Intersection

Definition

An oriented bounding box is a box whose faces have normals u , v , and w which make up an orthonormal basis. It is described by its center c and its positive half-lengths h_u , h_v , and h_w , i.e. the shortest distance from center to face along the respective normal directions.

Alternative Description


Consider a slab which is the infinitely large volume in between two parallel planes. For each direction consider the slab defined by the opposing faces of an OBB then the intersection of all three slabs equals the oriented bounding box.

Ray OBB Intersection

Slab Ray intersection

Each ray intersects the plane twice at (t_{u-}, t_{v-}, t_{w-}) and (t_{u+}, t_{v+}, t_{w+}) respectively. A ray hits the OBB if it enters all slabs before it leaves any

$$\max(t_{u-}, t_{v-}, t_{w-}) \leq \min(t_{u+}, t_{v+}, t_{w+}). \quad (3)$$



Ray OBB Intersection

```
1: procedure RaysOBBIntersect(Ray, OBB)
2:    $t_{\min} = -\infty$ 
3:    $t_{\max} = \infty$ 
4:   for direction  $a \in \{u, v, w\}$  do
5:      $e = a \cdot (c - o)$ 
6:      $f = a \cdot d$ 
7:     if  $|f| >$  floating point precision then
8:        $t_1 = (e + h_a)/f$ 
9:        $t_2 = (e - h_a)/f$ 
10:      if  $t_1 > t_2$  then swap( $t_1, t_2$ )
11:      end if
12:      if  $t_1 > t_{\min}$  then  $t_{\min} = t_1$ 
13:      end if
14:      if  $t_2 < t_{\max}$  then  $t_{\max} = t_2$ 
15:      end if
```

Ray OBB Intersection

```
16:         if  $t_{\min} > t_{\max}$  then return OBB not hit
17:         end if
18:         if  $t_{\max} < 0$  then return OBB behind camera
19:         end if
20:         else if  $-e - h_a > 0$  or  $-e + h_a < 0$  then return Ray
parallel to OBB but not hitting
21:         end if
22:     end for
23:     if  $t_{\min} > 0$  then return intersect at  $t_{\min}$ 
24:     else return intersect at  $t_{\max}$ 
25:     end if
26: end procedure
```

Ray Polygon Intersection

Definition Plane

A plane can be defined by a normal vector n and a point on the plane p by $\{x | n \cdot (x - p) = 0\}$. Alternatively, we can use the distance of the plane to the origin $\delta = n \cdot p$ to define the plane

$$\{x | n \cdot x - \delta = 0\}. \quad (4)$$

Definition Polygon

A simple planar polygon of n vertices (n -gon) is defined by an ordered vertex list $\{p_i\}_{i=1}^n$ and a plane with normal vector n and its distance to the origin δ . The vertices p_i and p_{i+1} form an edge for all $i = 1, \dots, n - 1$ and the polygon is closed by the edge from p_n to p_1 , where non of the edges intersect each other.

Ray Polygon Intersection

- To intersect a ray $p(t) = o + td$ and a polygon first calculate the intersection between the polygon plane and the ray

$$n \cdot (o + td) + \delta = 0. \quad (5)$$


- If $|n \cdot d| \ll 1$ the ray is considered parallel to the plane (avoid overflow when dividing), otherwise the intersection occurs for

$$t_{\text{intersect}} = \frac{-\delta - n \cdot o}{n \cdot d}. \quad (6)$$

- Project all vertices and the intersection point $p(t_{\text{intersect}})$ to one of the xy -, yz -, or xz -plane, where the area of the projected polygon is maximized and relabel the new coordinates by u and v . I.e. skip the coordinate component, where n is largest (skip y if $n = (1, 2, 0.5)$).
- Shift the projected intersection point into the origin.

Ray Polygon Intersection

- Count crossings with positive u -axis.
 - odd number: Origin inside polygon
 - even number: Origin outside polygon
 - infinite crossings: Rotate vertices slightly and repeat test



Ray Polygon Intersection



```
1: procedure CountCrossings
2:    $n_{\text{crossings}} = 0$ 
3:   for edge vertex pairs (a, b) do
4:     if ( $a_v > 0$  and  $b_v < 0$ ) or ( $a_v < 0$  and  $b_v > 0$ ) then
5:       if ( $a_u > 0$  and  $b_u > 0$ ) then
6:          $n_{\text{crossings}} = n_{\text{crossings}} + 1$ 
7:       else if intersection edge  $u$ -axis positive then
8:          $n_{\text{crossings}} = n_{\text{crossings}} + 1$ 
9:       else no crossing possible
10:      end if
11:    else no crossing possible
12:    end if
13:  end for return  $n_{\text{crossings}}$ 
14: end procedure
```

Instancing

Geometry instancing is the practice of rendering multiple copies of the same possibly distorted object in a scene. This method can also be applied to ray tracing where one can chose to

1. Transform the object and intersect with the untransformed ray
 2. Apply the inverse transformation to the ray and intersect with the untransformed object
- Homogeneous coordinates can be used to transform the ray origin (point) and ray direction.
 - Untransformed objects may have simpler intersection routines (sphere vs. ellipsoid)
 - Many transformed objects might share the same untransformed object reducing storage

Instancing



Optimization Techniques in Intersection

- Order comparisons such that easy to compute reject or accept intersections come first to possibly escape expensive computations
- Exploit the results from previous tests whenever possible
- Try out different orders if you have multiple tests. Minor changes might result in a significant performance boost
- Postpone expensive calculations such as trigonometric functions and square roots until they are required
- Dimensionality reduction simplify an intersection problem quite significantly
- In some cases pre-calculations can be worth the effort, though they require memory
- Take all special cases into account even if they are unlikely to occur
- Intersect with bounding boxes first, for complicated calculation

Bounding Volumes

- A Bounding Volume is a volume enclosing an object which is used for performant intersection rejection
- Common bounding volumes are spheres, axis aligned bounding boxes and oriented bounding boxes
- Finding optimal boxes is usually a non-trivial problem
- Often suboptimal boxes are chosen trading speed vs. efficiency
- In order to minimize intersection costs it is important to find tight fittings for an object.



Bounding Volumes

Axis Aligned Bounding Box (AABB)

- Find extreme coordinate values for each direction
- create tight axis aligned box from these values

Sphere

- Sphere center set to coincide with center from AABB
- radius given by largest distance to center position



Spatial Data Structures

- In case not only a few but a large number n of objects is considered the intersection costs increase by $\mathcal{O}(n)$
- Spatial data structures organizes objects in some space to accelerate intersection queries typically to $\mathcal{O}(\log n)$
- Most spatial data structures organize objects hierarchically
- Construction of these structures is expensive and usually done in a preprocessing step
- Common types are bounding volume hierarchies (BVHs) and binary space partitioning (BSP) trees
 - In BVHs each object belongs to exactly one of a number of nodes within each hierarchy level, whereas spatial points may belong to a number of nodes
 - In BSPs each spatial point belongs to exactly one of a number of nodes within each hierarchy level, whereas an object may belong to a number of nodes

Tree Structures in Computer Graphics

The scene is organized in a hierarchical tree structure consisting of a root, internal nodes, and leaves.

- The topmost node is the root containing all scene objects
- A leaf node holds a single scene object
- An internal node holds a number of children (nodes and/or leaves)
- Each node/leaf has a bounding volume enclosing all objects in its entire sub-tree




Figure 1: Example of bounding volume hierarchy (BVH) in two dimensions, where bounding volumes are AABB by Schreiberx licencend under CC BY-SA 3.0.

Find first object in hierarchy

- Start intersection with the root bounding volume
- If the ray misses it misses all objects
- otherwise continue recursively by testing the bounding volumes of the children
- Whenever a bounding volume is missed the entire sub-tree can be dismissed
- If the bounding volume of a leaf node is hit intersection with the object contained within is tested

Find closest object in hierarchy

- Store the identity and distance of the first object hit
- When intersecting bounding volumes discard sub-trees if the distance to the volume is larger than the stored distance
- Update identity and distance if a closer object is hit
- Continue until the whole tree is traversed

Bounding Volume Hierarchy

Example Hierarchy Creation

- 1: Chose an integer constant c
- 2: Decide on method to chose direction e , e.g. cycle through axis-directions as we go deeper into the hierarchy
- 3: Create AABB for all objects of scene
- 4: **procedure** bvh(objects)
- 5: **if** Number of objects $\leq c$ **then**
- 6: Create Bounding box enclosing all objects
- 7: Create a leaf for each object
- 8: **else**
- 9: Order objects along e (AABB center)
- 10: Partition objects into c sets
- 11: Call bhv for each set
- 12: **end if**
- 13: **end procedure**

Binary Space partitioning trees

Example Space Partitioning

- 1: Create AABB for all objects of scene
- 2: **procedure** bps(AABB, objects)
- 3: **if** Some continuation criterion is met **then**
- 4: Divide the AABB in two along an axis, e.g. the axis where
the AABB is largest
- 5: **for** each sub-box $AABB_i$ **do**
- 6: $O_i =$ set of objects (partially) inside $AABB_i$
- 7: bps($AABB_i$, O_i)
- 8: **end for**
- 9: **else**
- 10: create one leaf containing the objects
- 11: **end if**
- 12: **end procedure**

Remarks

- There is no recipe on which data structure to use
 - Depends on task to be accelerated
 - Depends on scene to be rendered
- There are more spatial data structures available, which might be useful depending on your application
 - Regular grids
 - Octrees
 - Constructive Solid Geometry trees