Ray tracing and geometric representations



http://graphics.cs.cmu.edu/courses/15-468

15-468, 15-668, 15-868 Physics-based Rendering Spring 2022, Lecture 2

Course announcements

- Programming assignment 0 is available on Canvas.
- Programming assignment 1 will be posted on Friday 1/28 and will be due two weeks later.
- Take-home quiz 1 will be posted on Tuesday 1/25 and will be due a week later.
- Office hours for this week only (will finalize starting next week based on survey results): \bullet - Yannis–Thursday 2-4 pm.
 - Zoom details on Piazza and Canvas.



2

Course announcements

Is anyone not on Piazza? \bullet

https://piazza.com/class/ky96bnus9u54ul

Is anyone not on Canvas? \bullet

Is anyone not on Slack?

https://canvas.cmu.edu/courses/27795

3

Overview of today's lecture

- Introduction to ray tracing.
- Intersections with geometric primitives.
- Triangular meshes.



Slide credits

Most of these slides were directly adapted from:

• Wojciech Jarosz (Dartmouth).

5

Two forms of 3D rendering

Rasterization: object point to image plane

- start with a 3D object point
- apply transforms
- determine the 2D image plane point it projects to

Ray tracing: image plane to object point

- start with a 2D image point
- generate a ray
- determine the visible 3D object point

Inverse processes



Two forms of 3D rendering

Rasterization

for (each triangle) for (each pixel or ray) for (each pixel) → for (each triangle) if (ray hits triangle) if (triangle covers pixel) keep closest hit keep closest hit **Triangle-centric Ray-centric**

Ray tracing





Rasterization advantages

Modern scenes are more complicated than images

- (not that much)
 - of course, if we have more than one sample per pixel this gets larger, but e.g. 4x supersampling is still a relatively comfortable ~100 MB
- Our scenes are routinely larger than this
 - This wasn't always true
- A rasterization-based renderer can *stream* over the triangles, no need to keep entire dataset around
- Allows parallelism and optimizations of memory systems

- A 1920x1080 frame (1080p) at 64-bit color and 32-bit depth per pixel is 24 MB



8

Rasterization limitations

Restricted to scan-convertible primitives

- Pretty much: triangles

Faceting, shading artifacts

- This is largely going away with programmable per-pixel shading, though

No unified handling of shadows, reflection, transparency







Ray/path tracing

Advantages

- Generality: can render anything that can be intersected with a ray - Easily allows recursion (shadows, reflections, etc.)
- Disadvantages
- Hard to implement in hardware (lacks computation coherence, must fit entire scene in memory, bad memory behavior)
 - Not such a big point anymore given general purpose GPUs
- Has traditionally been too slow for interactive applications
- Both of the above are changing rather rapidly right now!



A ray-traced image



11



Ray tracing today



Rapid change in film industry

2008:

- Most CGI in films rendered using micro-polygon rasterization.
- "You'd be crazy to render a full-feature film with ray/path tracing."
- Ray/path tracing mostly interesting to academics

2018:

- Most major films now rendered using ray/path tracing.
- "You'd be crazy *not* to render a full-feature film using path tracing."



Albrecht Dürer (1525)





René Descartes (1650)



15

Isaac Newton (1670)







Appel (1968)



Ray casting

- Generate an image by sending one ray per pixel
- Check for shadows by sending a ray towards the light



Whitted (1979)



recursive ray tracing (reflection & refraction)



Light Transport - Assumptions

Geometric optics:

- no diffraction, no polarization, no interference
- Light travels in a straight line in a vacuum
- no atmospheric scattering or refraction
- no gravity effects

Color can be represented as three numbers: (R,G,B)



Emission theory of vision



Eyes send out "feeling rays" into the world

- Supported by:
- Ancient greeks
- 50% of US college students*







Ray Tracing - Overview

"light tracing"





Ray Generation

























































Ray Tracing Pseudocode

rayTraceImage()

parse scene description

for each pixel ray = generateCameraRay(pixel) pixelColor = trace(ray)



Ray Tracing Pseudocode trace(ray) hit = find first intersection with scene objects color = (shade(hit)) return color might trace more rays (recursive)



Ray Tracing Pseudocode

rayTraceImage()

parse scene description

for each pixel ray = (generateCameraRay(pixel)) pixelColor = trace(ray) what is a ray?

how do we generate a camera ray?



Ray: a half line

Standard representation: origin (point) o and direction d

- this is a parametric equation for the line
- lets us directly generate the points on the line
- if we restrict to t > 0 then we have a ray
- note replacing d with ad does not change ray (for a > 0)

$\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$





Generating eye rays

Orthographic







Pinhole Camera (Camera Obscura)





Pinhole Camera





Pinhole Camera

Pinhole Camera




Generating eye rays—perspective

Establish view rectangle in X–Y plane, specified by, e.g. - l, r, t, b

Place rectangle at z = -d

$$\mathbf{s} = [u, v, -d]^T$$
$$\mathbf{d} = \mathbf{s}$$
$$\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$$

Does distance d matter?

 $\vee \geq t$



Placing the camera in the scene







Generating eye rays—orthographic

How do you generate a ray for an orthographic camera?









Surface primitives

- spheres
- planes
- triangles
- general implicits
- etc.



Ray-Sphere Intersection

Algebraic approach:

- Condition 1: point is on ray:
- Condition 2: point is on sphere:

- substitute and solve for *t*:

$$\|\mathbf{o} + t\mathbf{d} - \mathbf{c}\|^2 - r^2 = 0$$







Ray-Sphere Intersection

substitute and solve for t $\|\mathbf{o} + t\mathbf{d} - \mathbf{c}\|^2 - r^2 = 0 \longrightarrow (\mathbf{o}_x + t\mathbf{d}_x - \mathbf{c}_x)^2 +$

which reduces to: $At^2 + Bt + C = 0$ Solve for *t* using quadratic equation:

$$t = \frac{-B \pm \sqrt{B^2 - 4AC}}{2A}$$

What happens when square root is zero or negative?







Surface primitives

- spheres
- planes
- triangles
- general implicits
- etc.



Ray-Plane Intersection Plane equation (implicit)

Algebraic form:

ax + by + cz + d = 0











Surface primitives

- spheres
- planes
- triangles
- general implicits
- etc.



Ray-Triangle intersection Condition 1: point is on ray: Condition 2: point is on plane Condition 3: point is on the inside of all three edges First solve 1&2 (ray–plane intersection) for t: $(\mathbf{o} + t\mathbf{d})$ $t = -\frac{(\mathbf{o} - \mathbf{p}) \cdot \mathbf{n}}{\mathbf{d} \cdot \mathbf{n}}$

Several options for 3

$$\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$$

$$(\mathbf{x} - \mathbf{p}) \cdot \mathbf{n} = 0$$

$$-\mathbf{p}) \cdot \mathbf{n} = 0$$
$$\mathbf{p} - \mathbf{p} \cdot \mathbf{n}$$



In plane, triangle is the intersection of 3 half spaces





$\mathbf{n} = (\mathbf{p}_2 - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)$

Which way does n point?





$\mathbf{n} = (\mathbf{p}_2 - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)$ $\mathbf{n}_{\mathbf{x}13} = (\mathbf{x} - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)$ Which way does n point?

What about nx13?





$\mathbf{n} = (\mathbf{p}_2 - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)$ $\mathbf{n}_{x13} = (\mathbf{x} - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)$

- Which way does n point?
- What about n_{x13} ?
- How about now?





$\mathbf{n} = (\mathbf{p}_2 - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)$ $\mathbf{n}_{\mathbf{x}13} = (\mathbf{x} - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)$

- Which way does n point?
- What about n_{x13}?
- How about now?
- Edge test: $(\mathbf{n}_{\mathbf{x}13} \cdot \mathbf{n}) < 0$







$\mathbf{n} = (\mathbf{p}_2 - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)$ $\mathbf{n}_{\mathbf{x}13} = (\mathbf{x} - \mathbf{p}_1) \times (\mathbf{p}_3 - \mathbf{p}_1)$

- Which way does n point?
- What about nx13?
- How about now?
- Edge test: $(\mathbf{n}_{\mathbf{x}13} \cdot \mathbf{n}) < 0$







Intersect ray with triangle's plane

Test whether hit-point is within triangle

- compute sub-triangle areas α, β, γ
- test inside triangle conditions



Barycentric coordinates

Barycentric coordinates:

Inside triangle conditions:



 $\mathbf{x}(\alpha, \beta, \gamma) = \alpha \mathbf{p}_1 + \beta \mathbf{p}_2 + \gamma \mathbf{p}_3$ $\alpha + \beta + \gamma = 1 \quad 0 \le \alpha \le 1$ $\gamma = 1 - \alpha - \beta \quad 0 \le \beta \le 1$ $0 \le \gamma \le 1$



Interpretations of barycentric coords

Sub-triangle areas



 $\alpha = |\Delta \mathbf{p}_2 \mathbf{p}_3 \mathbf{x}| / |\Delta \mathbf{p}_1 \mathbf{p}_2 \mathbf{p}_3|$ $\beta = |\Delta \mathbf{p}_1 \mathbf{p}_3 \mathbf{x}| / |\Delta \mathbf{p}_1 \mathbf{p}_2 \mathbf{p}_3|$ $\gamma = |\Delta \mathbf{p}_1 \mathbf{p}_2 \mathbf{x}| / |\Delta \mathbf{p}_1 \mathbf{p}_2 \mathbf{p}_3|$

Ray-Triangle Intersection (Approach 3) $\alpha \mathbf{p}_1 + \beta \mathbf{p}_2 + (1 - \alpha - \beta)\mathbf{p}_3 = \mathbf{o} + t\mathbf{d}$ Insert ray equation: $\alpha(\mathbf{p}_1 - \mathbf{p}_3) + \beta(\mathbf{p}_2 - \mathbf{p}_3) + \mathbf{p}_3 = \mathbf{o} + t\mathbf{d}$ $\alpha(\mathbf{p}_1 - \mathbf{p}_3) + \beta(\mathbf{p}_2 - \mathbf{p}_3) - t\mathbf{d} = \mathbf{o} - \mathbf{p}_3$ $\alpha \mathbf{a} + \beta \mathbf{b} - t\mathbf{d} = \mathbf{e}$ $\begin{bmatrix} -\mathbf{d} & \mathbf{a} & \mathbf{b} \end{bmatrix} \begin{bmatrix} t \\ \alpha \\ \beta \end{bmatrix} = \mathbf{e}$ Solve directly Can be much faster!





Other primitives

- cylinder
- cone, paraboloid, hyperboloid
- torus
- disk
- general polygons, meshes
- etc.



Intersecting transformed primitive?

Option 1: Transform the primitive

- simple for triangles, since they transform to triangles
- other primitives get more complicated (e.g. sphere \rightarrow ellipsoid)
- Option 2: Transform the ray (by the inverse transform)
- more elegant; works on any primitive
- allows simpler intersection tests (e.g., just use existing sphere-intersection routine)



World space



Local space





World space



Local space





World space



Local space





Transformations in homogeneous coords

 $\mathbf{v} = \begin{pmatrix} x \\ y \\ z \end{pmatrix}$

A 3D transformation matrix:

A 3D homogenous vector:

A position has $w \neq 0$, and a direction has w = 0

 $M = \begin{pmatrix} m_{11} & m_{12} & m_{13} & m_{14} \\ m_{21} & m_{22} & m_{24} & m_{24} \\ m_{31} & m_{32} & m_{33} & m_{34} \\ m_{41} & m_{42} & m_{43} & m_{44} \end{pmatrix}$



Transformations

Matrix-vector multiplication, Mv, transforms the vector





Have a transform M, a ray $\mathbf{r}(t)$, and a surface S

- To intersect:
- 1. Transform ray to local coords (by inverse of M)
- 2. Call surface intersection
- 3. Transform hit data back to global coords (by M)
- How to transform a ray $\mathbf{r}(t) = \mathbf{p} + t\mathbf{d}$ by M^{-1} ?
- $\mathbf{r'}(t) = M^{-1}\mathbf{p} + tM^{-1}\mathbf{d}$
- Remember: p forms as a point, d as a direction!



Other primitives

- cylinder
- cone, paraboloid, hyperboloid
- torus
- disk
- general polygons, meshes
- etc.



Image so far With eye ray generation and sphere intersection parse scene description for each pixel: ray = camera.getRay(pixel); hit = s.intersect(ray, 0, +inf); if hit: image.set(pixel, white);







Intersecting many shapes

Intersect each primitive

Pick closest intersection

- Only within considered range [t_{min}, t_{max}]
- After each valid intersection, update t_{max}
- Essentially a line search







Intersection against many shapes

The basic idea is:

Surfaces::intersect(ray, tMin, tMax): tBest = +inf; firstHit = null; for s in surfaces: hit = s.intersect(ray, tMin, tBest); if hit: tBest = hit.t; firstHit = hit; return firstHit;

(acceleration structures)

- this is linear in number of surfaces but there are sublinear methods



Image so far

With eye ray generation and scene intersection

```
for each pixel:
ray = camera.getRay(pixel);
c = scene.trace(ray, 0, +inf);
image.set(pixel, c);
```

```
Scene::trace(ray, tMin, tMax):
hit = surfaces.intersect(ray, tMin, tMax);
if (hit)
   return hit.color();
else
   return backgroundColor;
```





Other primitives

- cylinder
- cone, paraboloid, hyperboloid
- torus
- disk
- general polygons, meshes
- etc.


How should we represent complex geometry?

How are they obtained?

- modeled by hand
- scanned

What operations must we support?

- modeling/editing
- animating
- texturing
- rendering







Surface representation zoo!

Parametric





- Splines, tensor-product surfaces
- Subdivision surfaces





After a slide by Olga Sorkine-Hornung



Polygonal Meshes





A small triangle mesh



lide by Steve Marschner S After a

12 triangles, 8 vertices



A large mesh

10 million triangles from a highresolution 3D scan









After a slide by Steve Marschner







spheres



approximate sphere

Rineau & Yvinec

CGAL manual



Meshes as Approx. of Smooth Surfaces

Piecewise linear approximation - Error is $O(h^2)$

lide by Olga Sorkine-Hornung

S

After a





Meshes as Approx. of Smooth Surfaces **Piecewise linear approximation** - Error is $O(h^2)$ #faces vs. approximation error 25. 20. Approximation error 15. 3 10. 5. 0. 15 23 8 0 30 Number of faces 25%

lide by Olga Sorkine-Hornung S After a





Polygonal Meshes

Polygonal meshes are a good representation

- approximation $O(h^2)$
- arbitrary topology
- piecewise smooth surfaces
- adaptive refinement
- efficient rendering





Data Structures: What should be stored?



- Attributes
- Normal, color, texture coordinates
- Per vertex, face, edge
- Connectivity
- Adjacency relationships

Geometry: 3D coordinates



Separate Triangle List or Face Set (STL)

Face: 3 vertex positions

Storage:

- 4 Bytes/coordinate (using 32-bit floats)
- 36 Bytes/face

Wastes space

Triangles				
0	x0	уO		
1	x1	y1		
2	x2	y2		
3	х3	у3		
4	x4	y4		
5	x5	y5		
6	x6	y6		
•••	•••	•••		





Indexed Face Set (OBJ, OFF, WRL)

Vertex: position

Face: vertex indices

Storage:

- 12 Bytes/vertex
- 12 Bytes/face

Reduces wasted space

Even better with per-vertex attributes



Triangles				
tO	V0	v1	v2	
t1	V0	v1	v3	
t2	v2	v4	v3	
t3	v5	v2	v6	
• • •	• • •	• • •	• • •	







Data on meshes

Often need to store additional information besides just the geometry

Can store additional data at faces, vertices, or edges Examples

- colors stored on faces, for faceted objects
- information about sharp creases stored at edges
- any quantity that varies continuously (without sudden changes, or *discontinuities*) gets stored at vertices



Key types of vertex data

- Surface normals
- vertices
- Texture coordinates
- 2D coordinates that tell you how to paste images on the surface
- Positions
- at some level this is just another piece of data

- when a mesh is approximating a curved surface, store normals at



Defining normals

Face normals: same normal for all points in facegeometrically correct, but faceted look







Problems with face normals

Piecewise planar approximation converges pretty quickly to the smooth geometry as the number of triangles increases

- error is $O(h^2)$

But the surface normals don't converge so well

- normal is constant over each trian edges
- error is only O(h)

- normal is constant over each triangle, with discontinuous jumps across



Problems with face normals—2D example

Approximating circle with increasingly many segments

Max error in position error drops by factor of 4 each step

Max error in normal only drops by factor of 2



Problems with face normals—solution

Piecewise planar approximation converges pretty quickly to the smooth geometry as the number of triangles increases

- for mathematicians: error is $O(h^2)$
- But the surface normals don't converge so well
- normal is constant over each triangle, with discontinuous jumps across edges
- for mathematicians: error is only O(h)

normals that vary gradually across triangles

- Better: store the "real" normal at each vertex, and interpolate to get



Defining normals

Vertex normals: store normal at vertices, interpolate in face - geometrically "inconsistent", but smooth look







Barycentric coordinates

Barycentric interpolation:



 $\mathbf{p}(\alpha,\beta,\gamma) = \alpha \mathbf{p}_1 + \beta \mathbf{p}_2 + \gamma \mathbf{p}_3$

Can use this eqn. to interpolate any vertex quantity across triangle!



Barycentric coordinates

Barycentric interpolation:



$\mathbf{p}(\alpha, \beta, \gamma) = \alpha \mathbf{p}_1 + \beta \mathbf{p}_2 + \gamma \mathbf{p}_3$ $\mathbf{c}(\alpha, \beta, \gamma) = \alpha \mathbf{c}_1 + \beta \mathbf{c}_2 + \gamma \mathbf{c}_3$

Can use this eqn. to interpolate any vertex quantity across triangle!



Barycentric coordinates

Barycentric interpolation:



 $\mathbf{p}(\alpha,\beta,\gamma) = \alpha \mathbf{p}_1 + \beta \mathbf{p}_2 + \gamma \mathbf{p}_3$ $\mathbf{c}(\alpha,\beta,\gamma) = \alpha \mathbf{c}_1 + \beta \mathbf{c}_2 + \gamma \mathbf{c}_3$ $\mathbf{n}(\alpha,\beta,\gamma) = \alpha \mathbf{n}_1 + \beta \mathbf{n}_2 + \gamma \mathbf{n}_3$

not guaranteed to be unit length

Can use this eqn. to interpolate any vertex quantity across triangle!



Realism through geometric complexity

DIDARABASA

Andreas Byström

100000000000000



Ray Tracing Acceleration

Ray-surface intersection is at the core of every ray tracing algorithm

Brute force approach:

- intersect every ray with every primitive
- many unnecessary raysurface intersection tests





Ray Tracing Cost

over 95 percent" [Whitted 1980]

 $Cost = O(n_x \cdot n_y \cdot n_o)$

- (number of pixels) · (number of objects)
- Assumes 1 ray per pixel

Example: 1024 x 1024 image of a scene with 1000 triangles

- Cost is (at least) 10⁹ ray-triangle intersections

Typically measured per ray:

- Naive: $O(n_o)$ - linear with number of objects

"the time required to compute the intersections of rays and surfaces is





O(n_o) Ray Tracing (The Problem)



8 primitives \rightarrow 3 seconds

50K trees each with 1M polygons = 50B polygons

 \rightarrow 594 years!



Sub-linear Ray Tracing



50K trees each with 1M polygons = 50B polygons \rightarrow **11 minutes 300,000,000x speedup!**

The solution

Improve efficiency of ray-surf acceleration structures.

- A spatial organization of objects in a scene to minimize the necessary number of ray-object intersection tests.
- Spatial sorting/subdivision (e.g. grid, kd-tree, ochre)
- Decompose space into disjoint regions & assign objects to regions
- Object sorting/subdivision (bounding volume hierarchy)
- Decompose objects into disjoint sets & bound using simple volumes for fast rejection

Improve efficiency of ray-surface intersections by constructing





Bounding Volumes

Spheres





Bounding Volumes

Axis-aligned bounding boxes (most common)







Bounding Volumes Hierarchies

Now do this hierarchically!





BVH Traversal

void BVHNode::intersectBVH(ray, &hit): if (bound.hit(ray)): if (leaf): else:

- leaf.intersect(ray, hit);
- leftChild.intersectBVH(ray, hit); rightChild.intersectBVH(ray, hit);



Constructing BVHs

- Top-down:
- partition objects along an axis and create two sub-sets
- Bottom-up:
- recursively group nearby objects together



Divisive (top-down) BBH construction

- 1. Choose split axis
- 2. Choose split plane location
- 3. Choose whether to create leaf or split + repeat
- Many strategies for each of these steps

leaf or split + repeat nese steps


Choosing axis based on centroid extents



PBRe2 fig. 4.8



Object-median splitting

- 1. Sort bbox centroids along split axis
- 2. Take take first half as left child, second half as right



split axis hild, second half as right

