

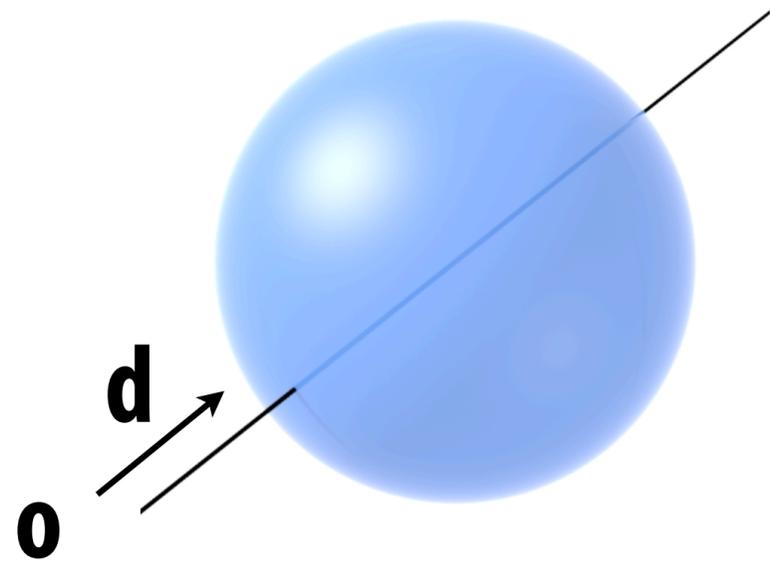
**Lecture 9:**

# **Accelerating Geometric Queries**

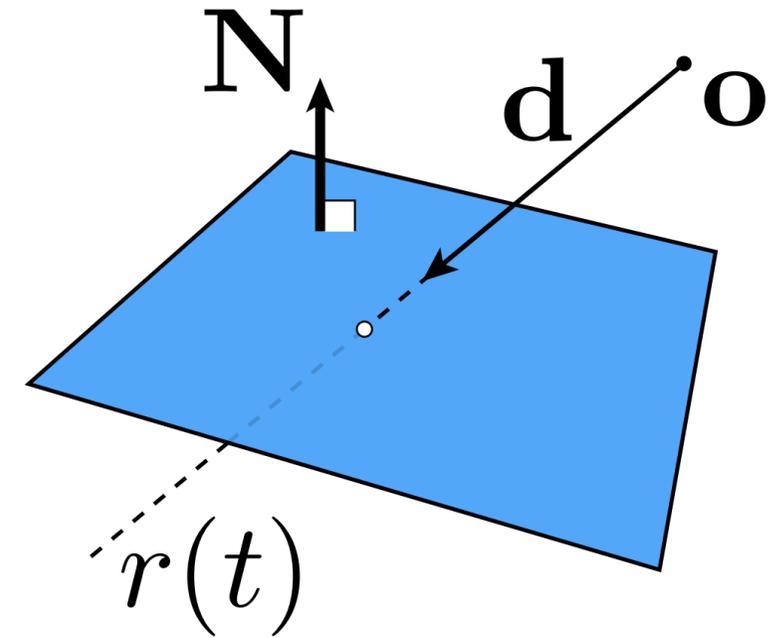
---

**Interactive Computer Graphics  
Stanford CS248, Spring 2018**

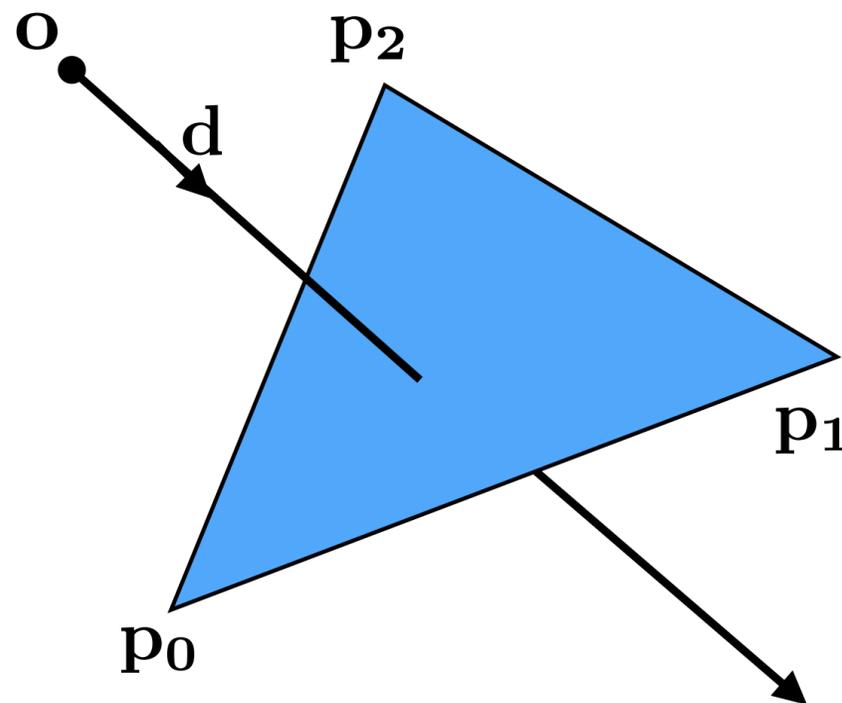
# Last time: intersecting a ray with individual primitives



Ray-sphere



Ray-plane



Ray-triangle

# Ray-scene intersection

Given a scene defined by a set of  $N$  primitives and a ray  $r$ , find the closest point of intersection of  $r$  with the scene

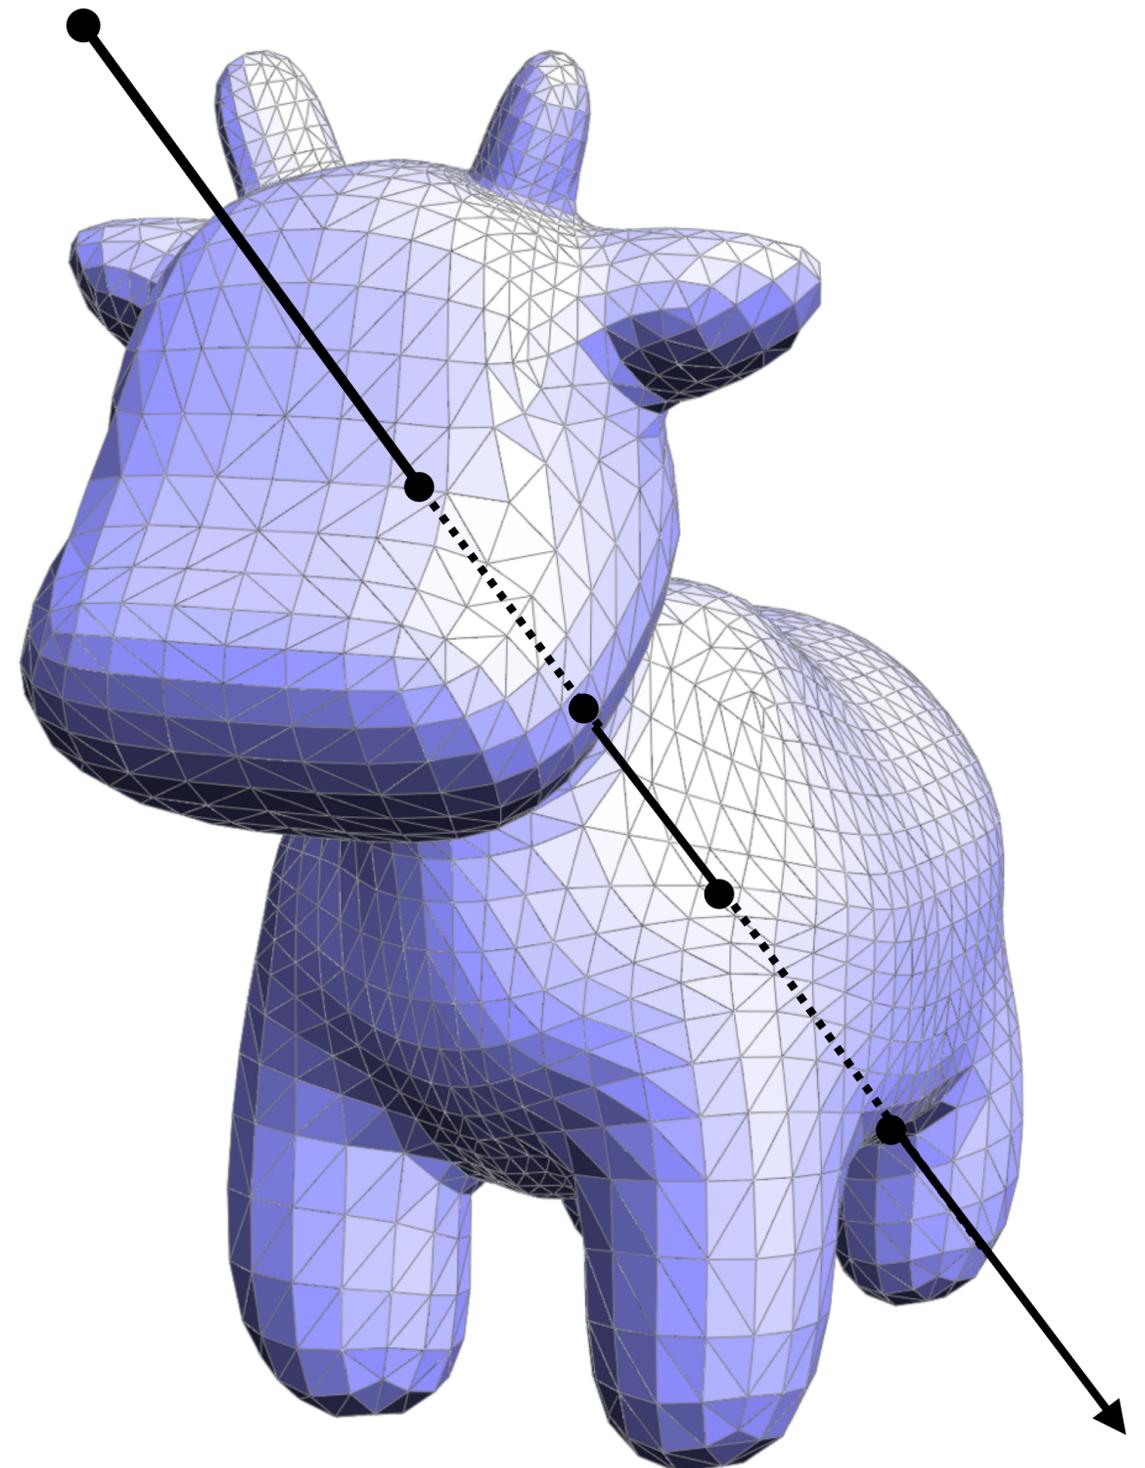
“Find the first primitive the ray hits”

```
p_closest = NULL
t_closest = inf
for each primitive p in scene:
    t = p.intersect(r)
    if t >= 0 && t < t_closest:
        t_closest = t
        p_closest = p
```

Complexity?  $O(N)$

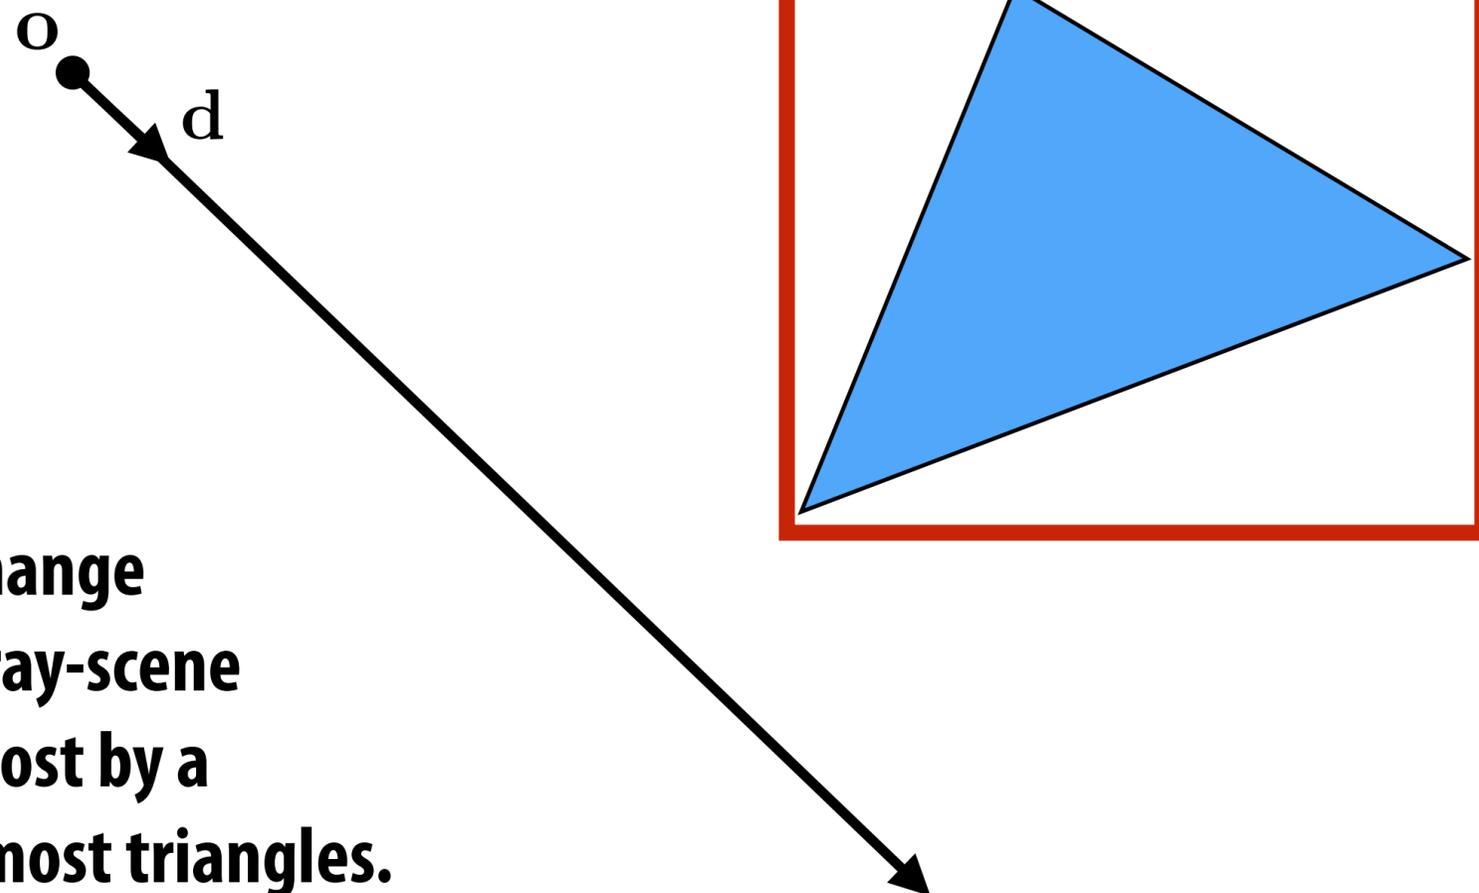
*Can we do better?*

(Assume `p.intersect(r)` returns value of  $t$  corresponding to the point of intersection with ray  $r$ )



# One simple idea

- **“Early out” — Skip ray-primitive test if it is computationally easy to determine that ray does not intersect primitives**
- **E.g., A ray cannot intersect a primitive if it doesn't intersect the bounding box containing it!**



**Note: early out does not change asymptotic complexity of ray-scene intersection. But reduces cost by a constant if ray is far from most triangles.**

# Ray-axis-aligned-box intersection

What is ray's closest/farthest intersection with axis-aligned box?

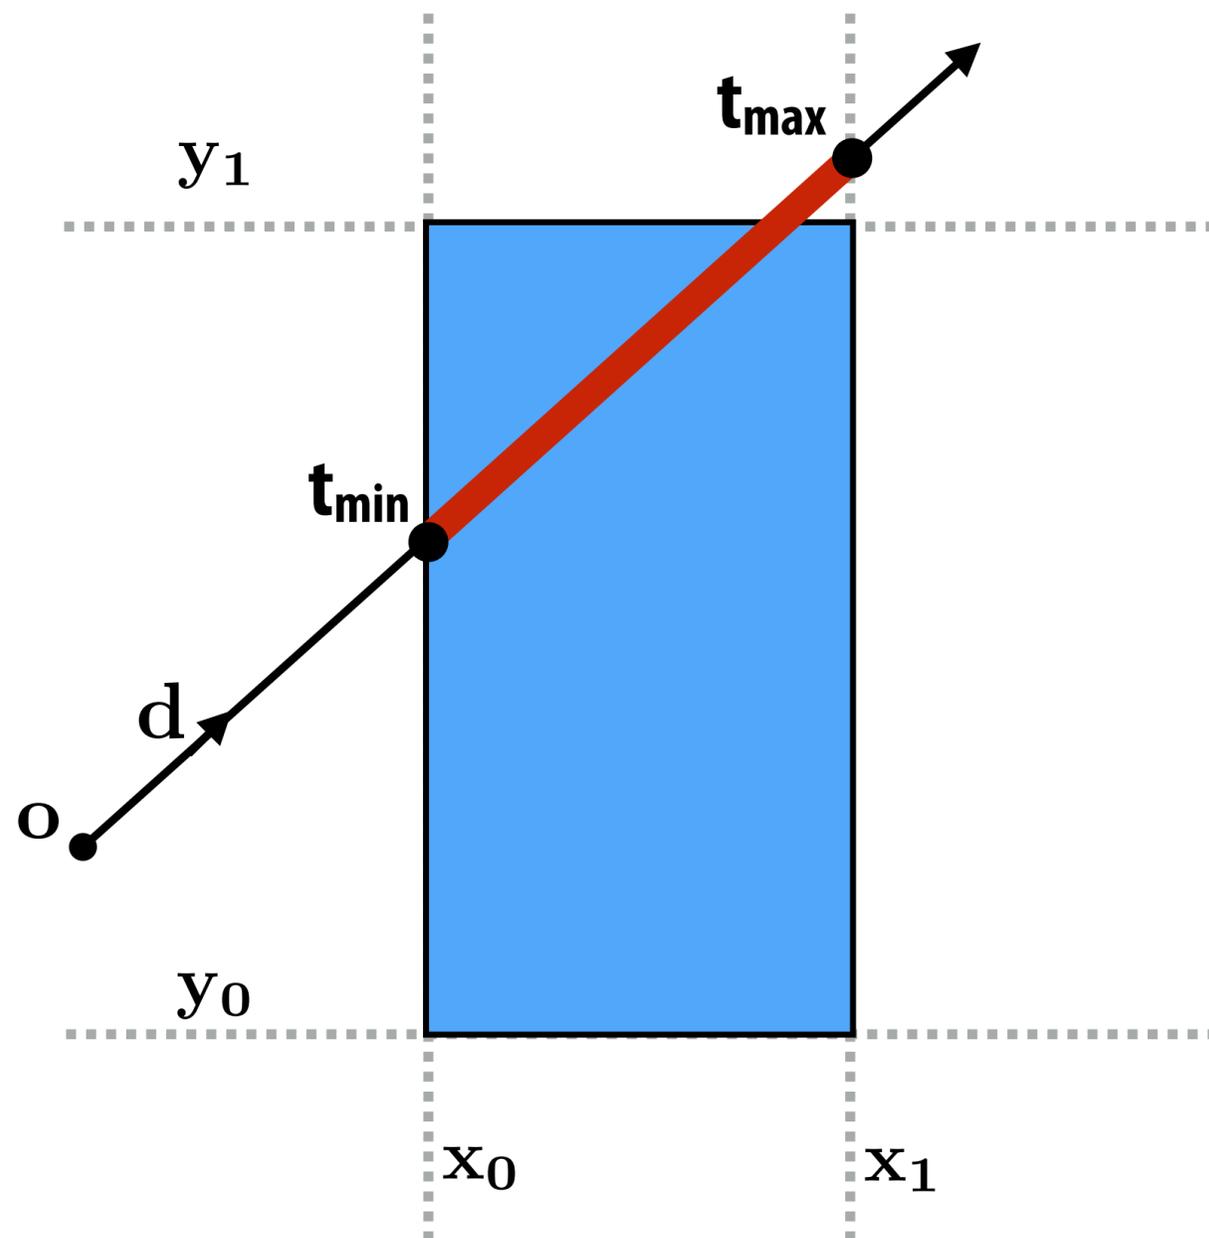


Figure shows intersections with  $x=x_0$  and  $x=x_1$  planes.

Find intersection of ray with all planes of box:

$$\mathbf{N}^T (\mathbf{o} + t\mathbf{d}) = c$$

Math simplifies greatly since plane is axis aligned (consider  $x=x_0$  plane in 2D):

$$\mathbf{N}^T = [1 \quad 0]^T$$

$$c = x_0$$

$$t = \frac{x_0 - \mathbf{O}_x}{d_x}$$

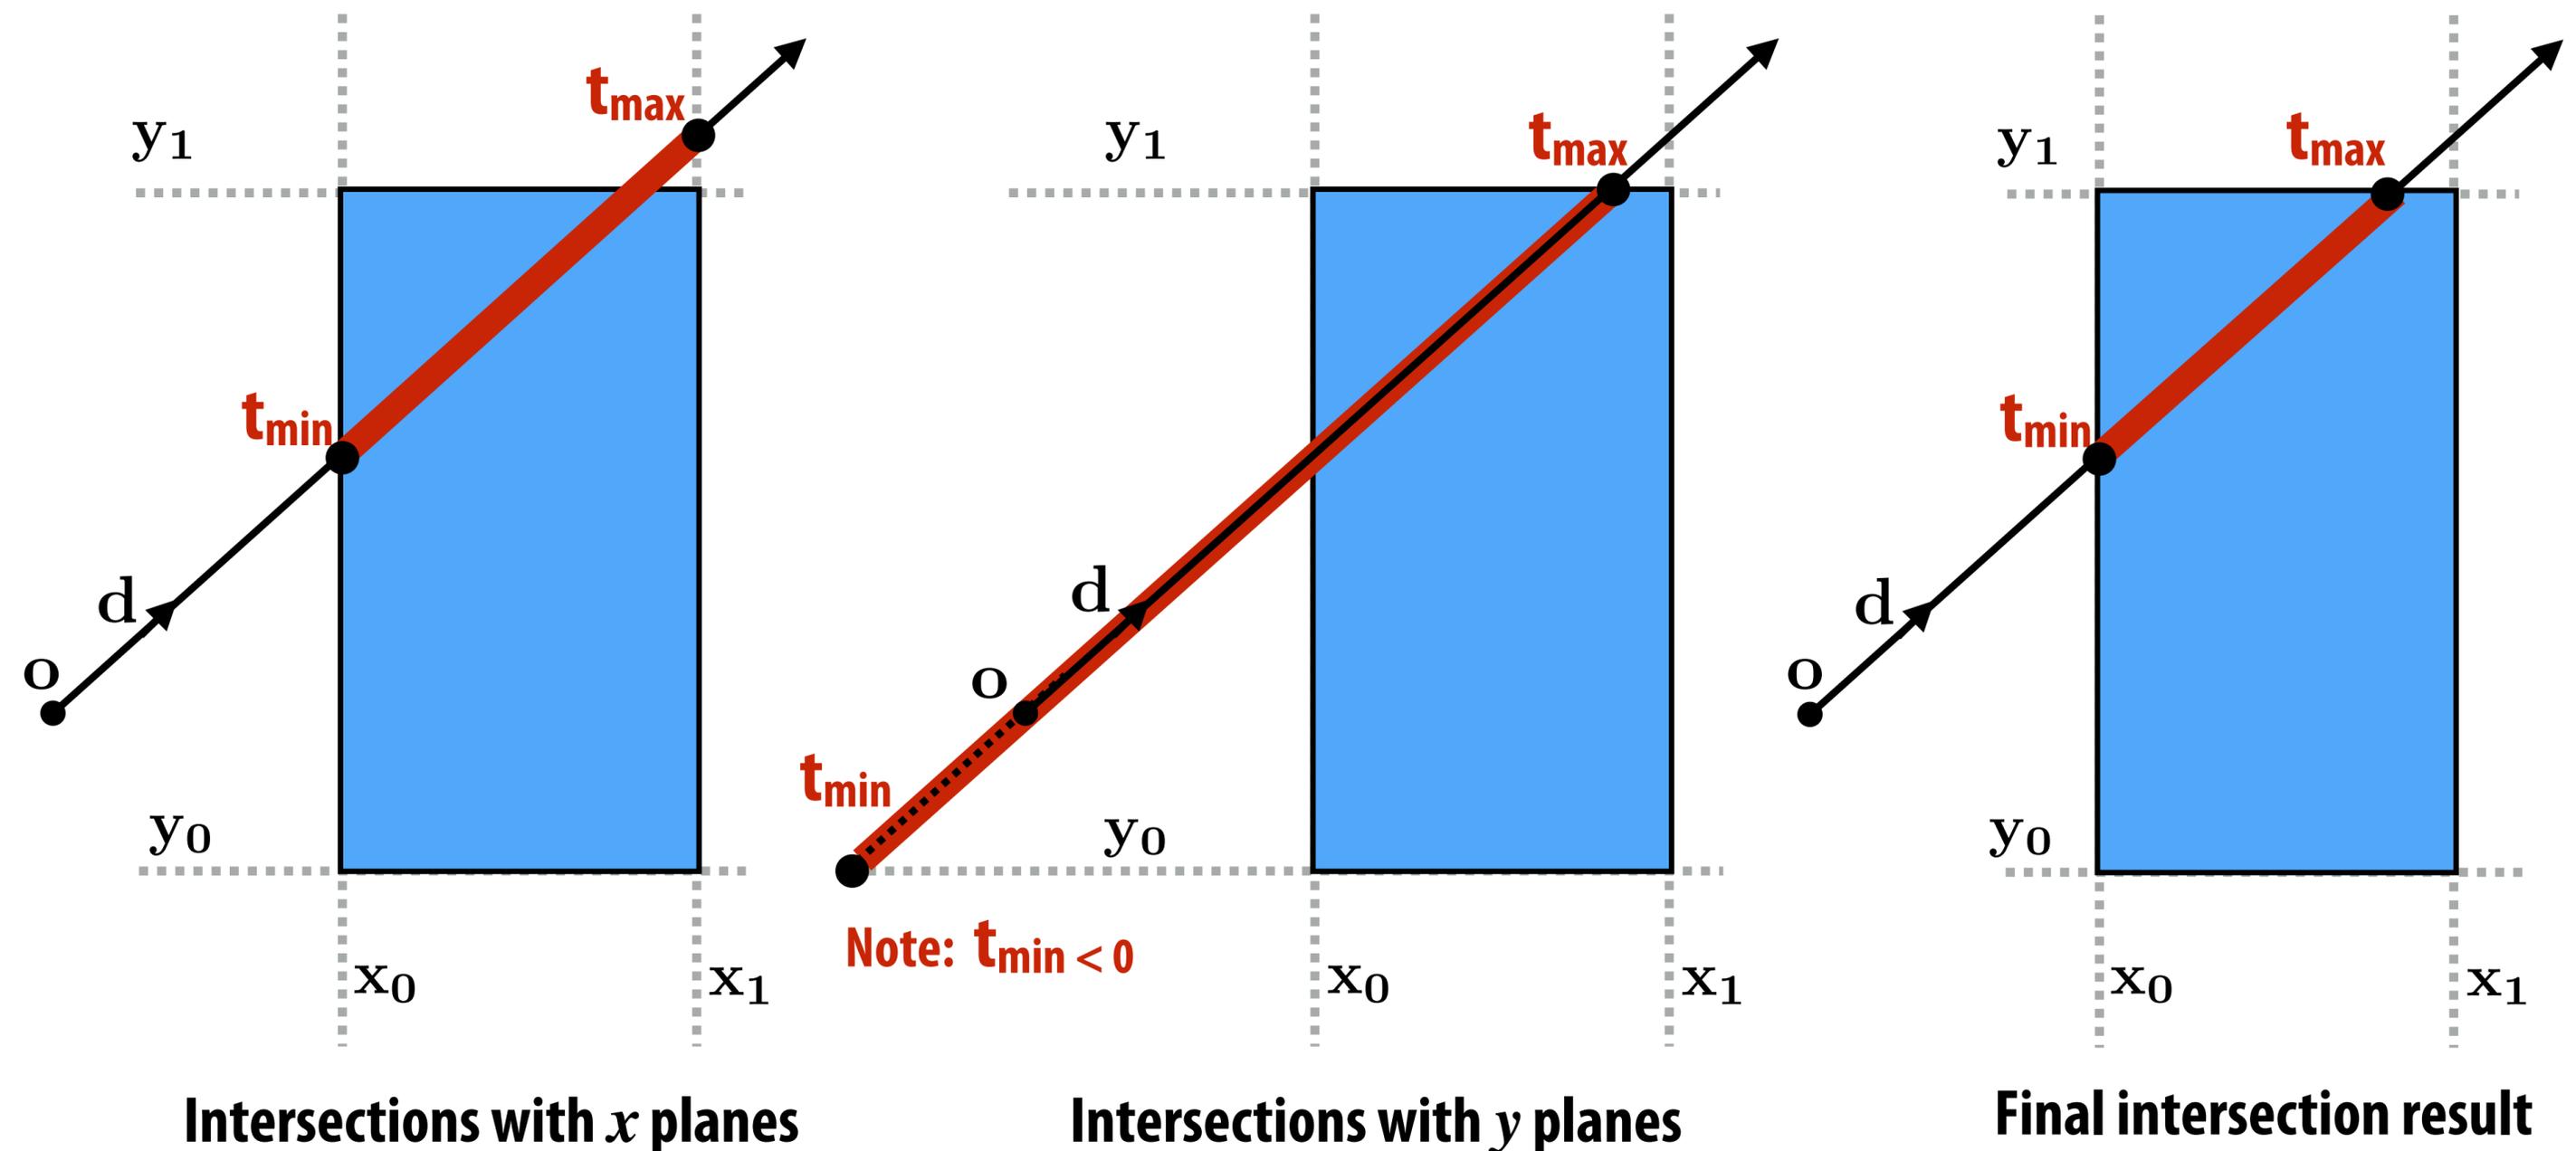
Performance note: it is possible to precompute box independent terms, so computing  $t$  is cheap

$$a = \frac{1}{d_x} \quad \text{and} \quad b = -\frac{\mathbf{O}_x}{d_x}$$

So...  $t = ax + b$

# Ray-axis-aligned-box intersection

Compute intersections with all planes, take intersection of  $t_{\min}/t_{\max}$  intervals



How do we know when the ray misses the box?

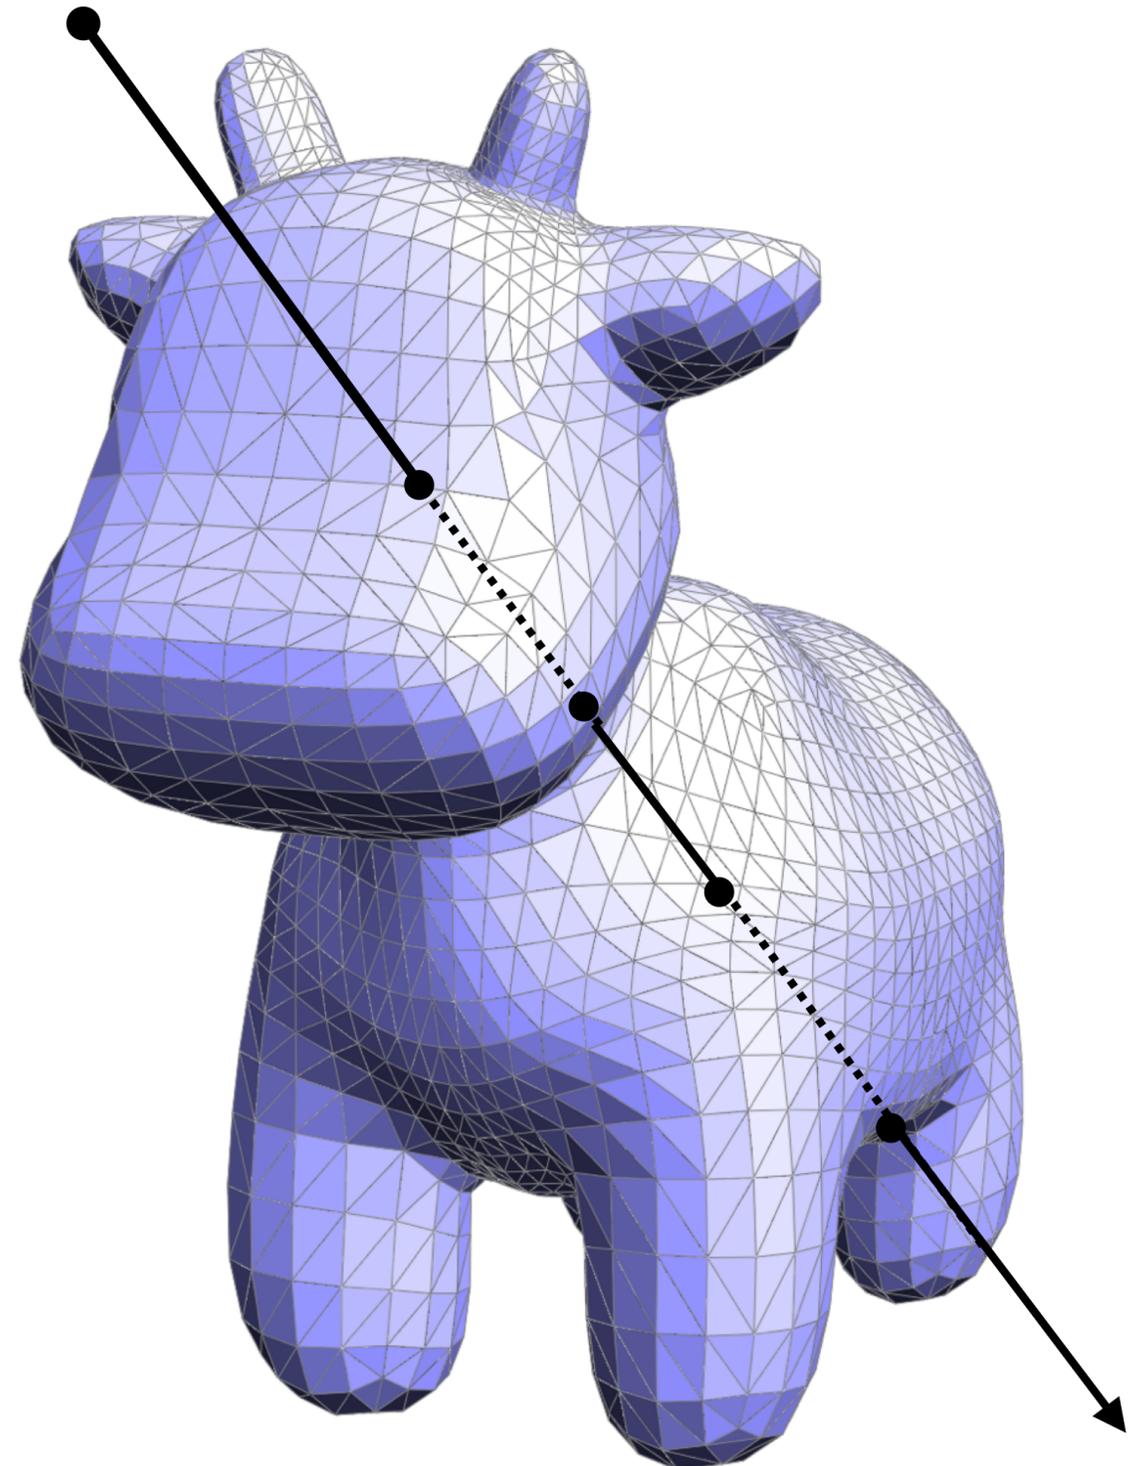
# Ray-scene intersection with early out

Given a scene defined by a set of  $N$  primitives and a ray  $r$ , find the closest point of intersection of  $r$  with the scene

```
p_closest = NULL
t_closest = inf
for each primitive p in scene:
    if (!p.bbox.intersect(r))
        continue;
    t = p.intersect(r)
    if t >= 0 && t < t_closest:
        t_closest = t
        p_closest = p
```

*Can we do better?*

(Assume `p.intersect(r)` returns value of  $t$  corresponding to the point of intersection with ray  $r$ )



# A simpler problem

- Imagine I have a set of integers  $S$
- Given an integer, say  $k=18$ , find the element of  $S$  closest to  $k$ :

10   123   2   100   6   25   64   11   200   30   950   111   20   8   1   80

What's the cost of finding  $k$  in terms of the size  $N$  of the set?

Can we do better?

Suppose we first *sort* the integers:

1   2   6   8   10   11   20   25   30   64   80   100   111   123   200   950

How much does it now cost to find  $k$  (*including sorting*)?

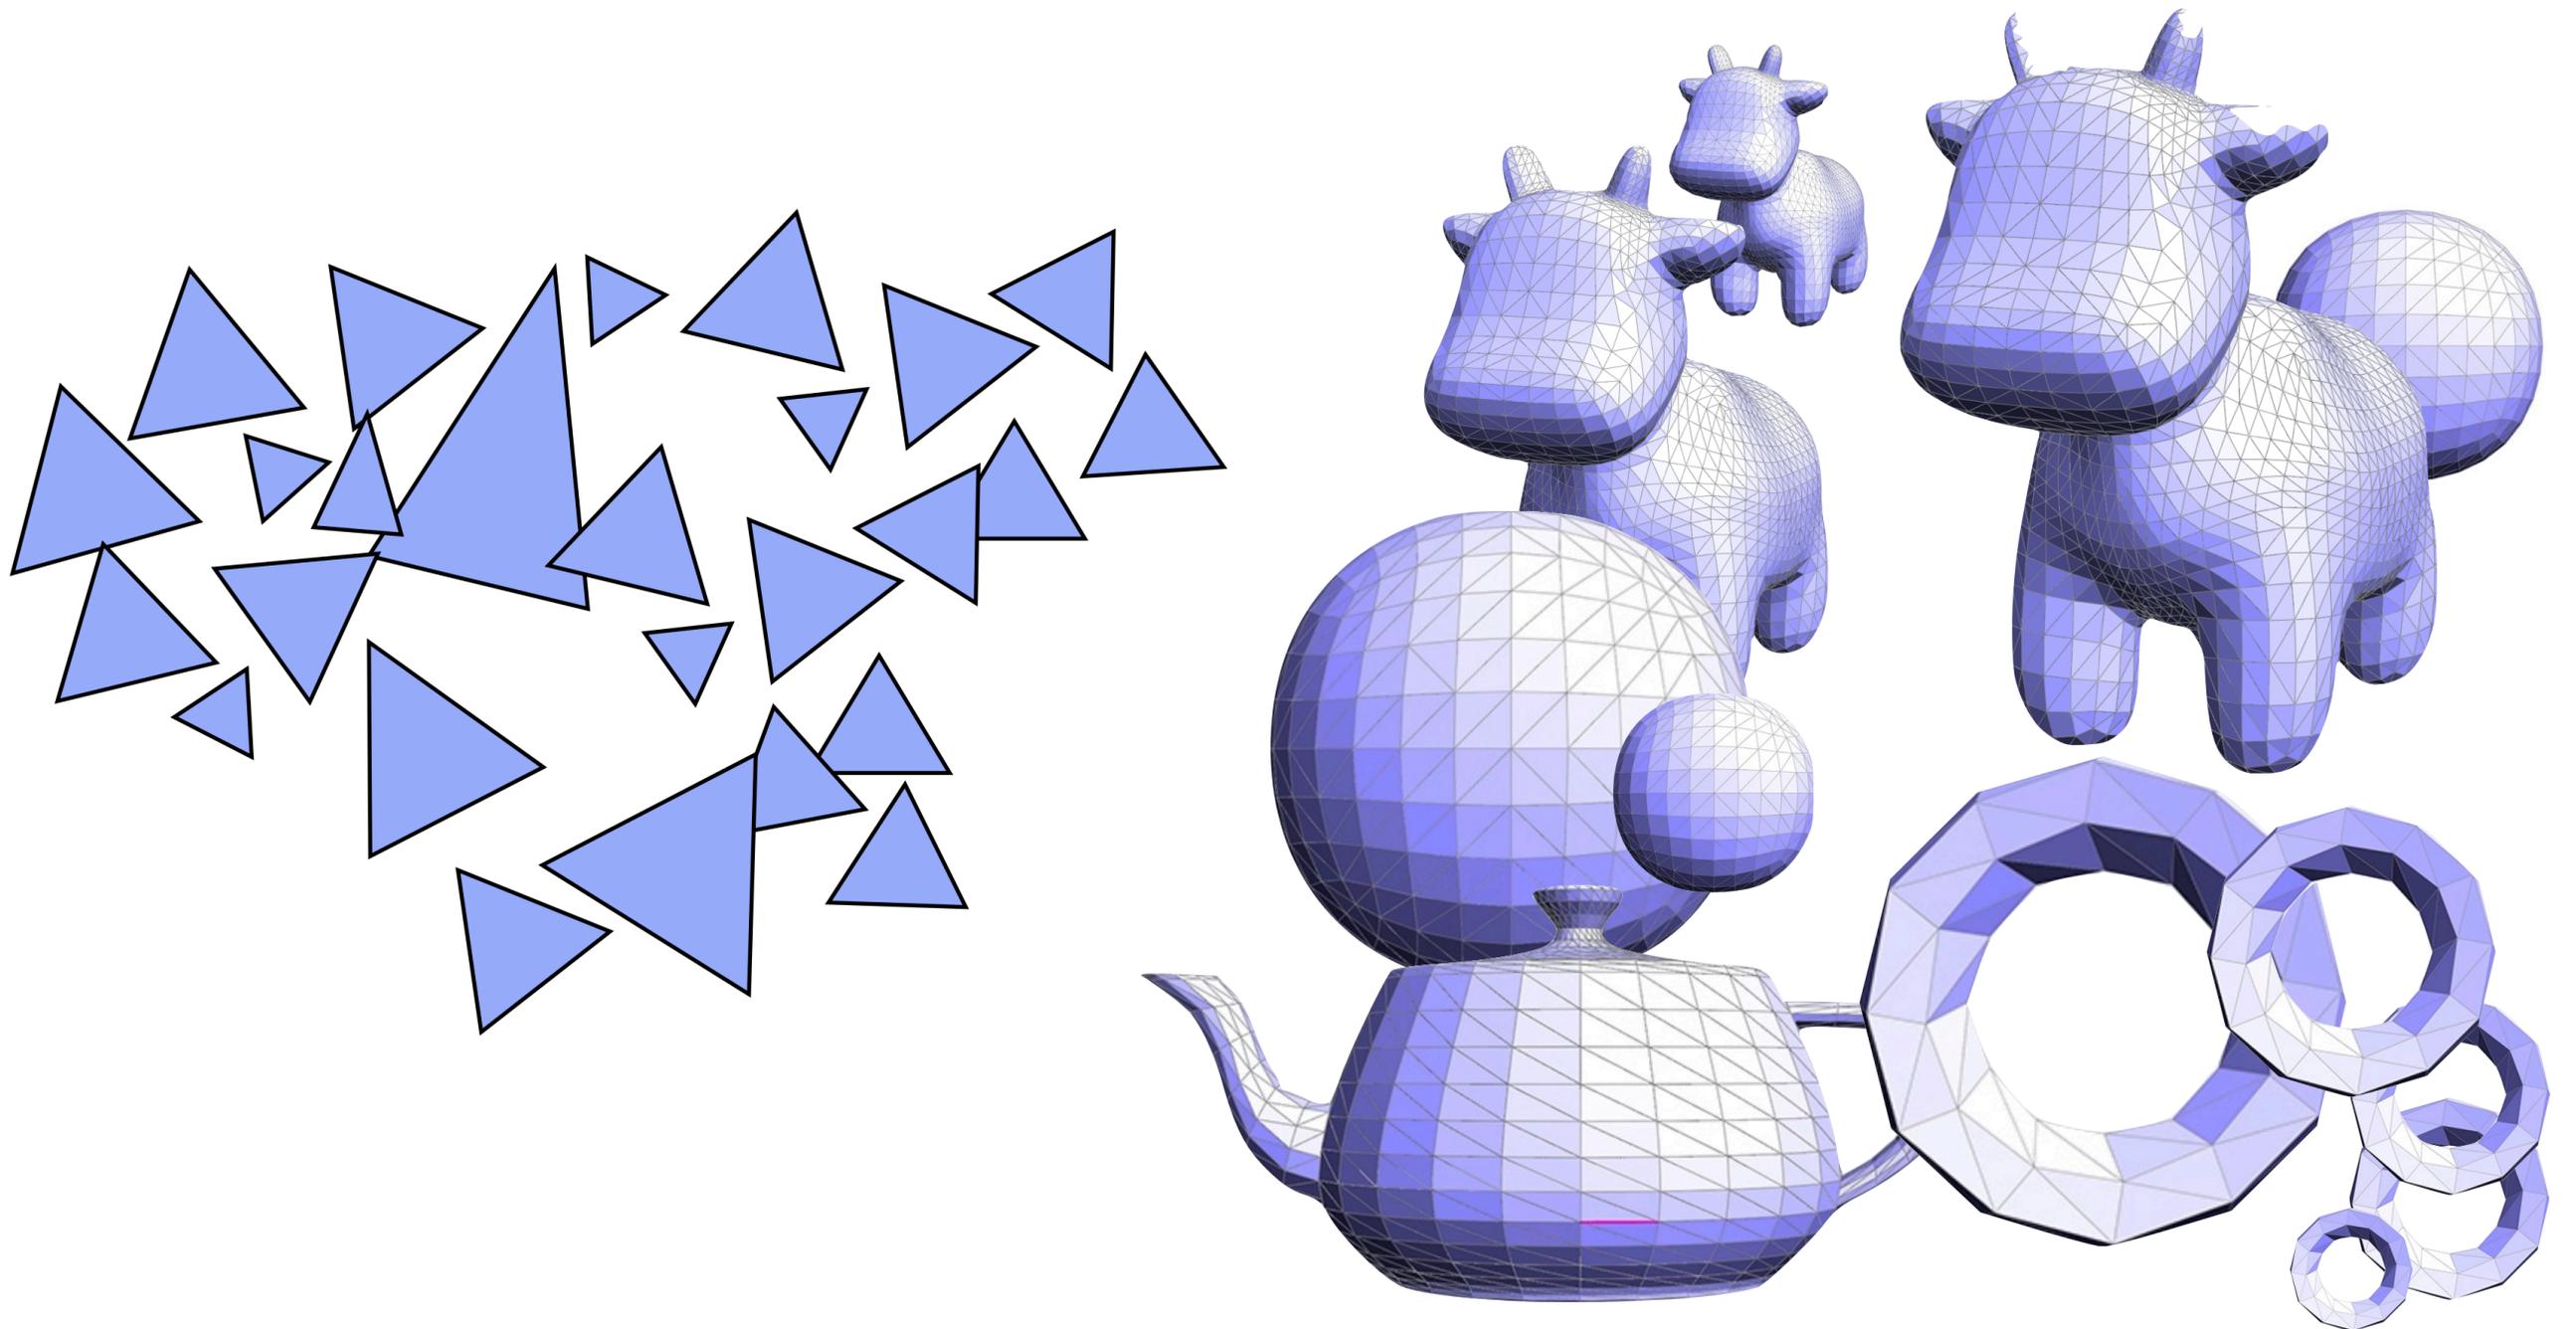
Cost for just ONE query:  $O(n \log n)$

Amortized cost over many queries:  $O(\log n)$

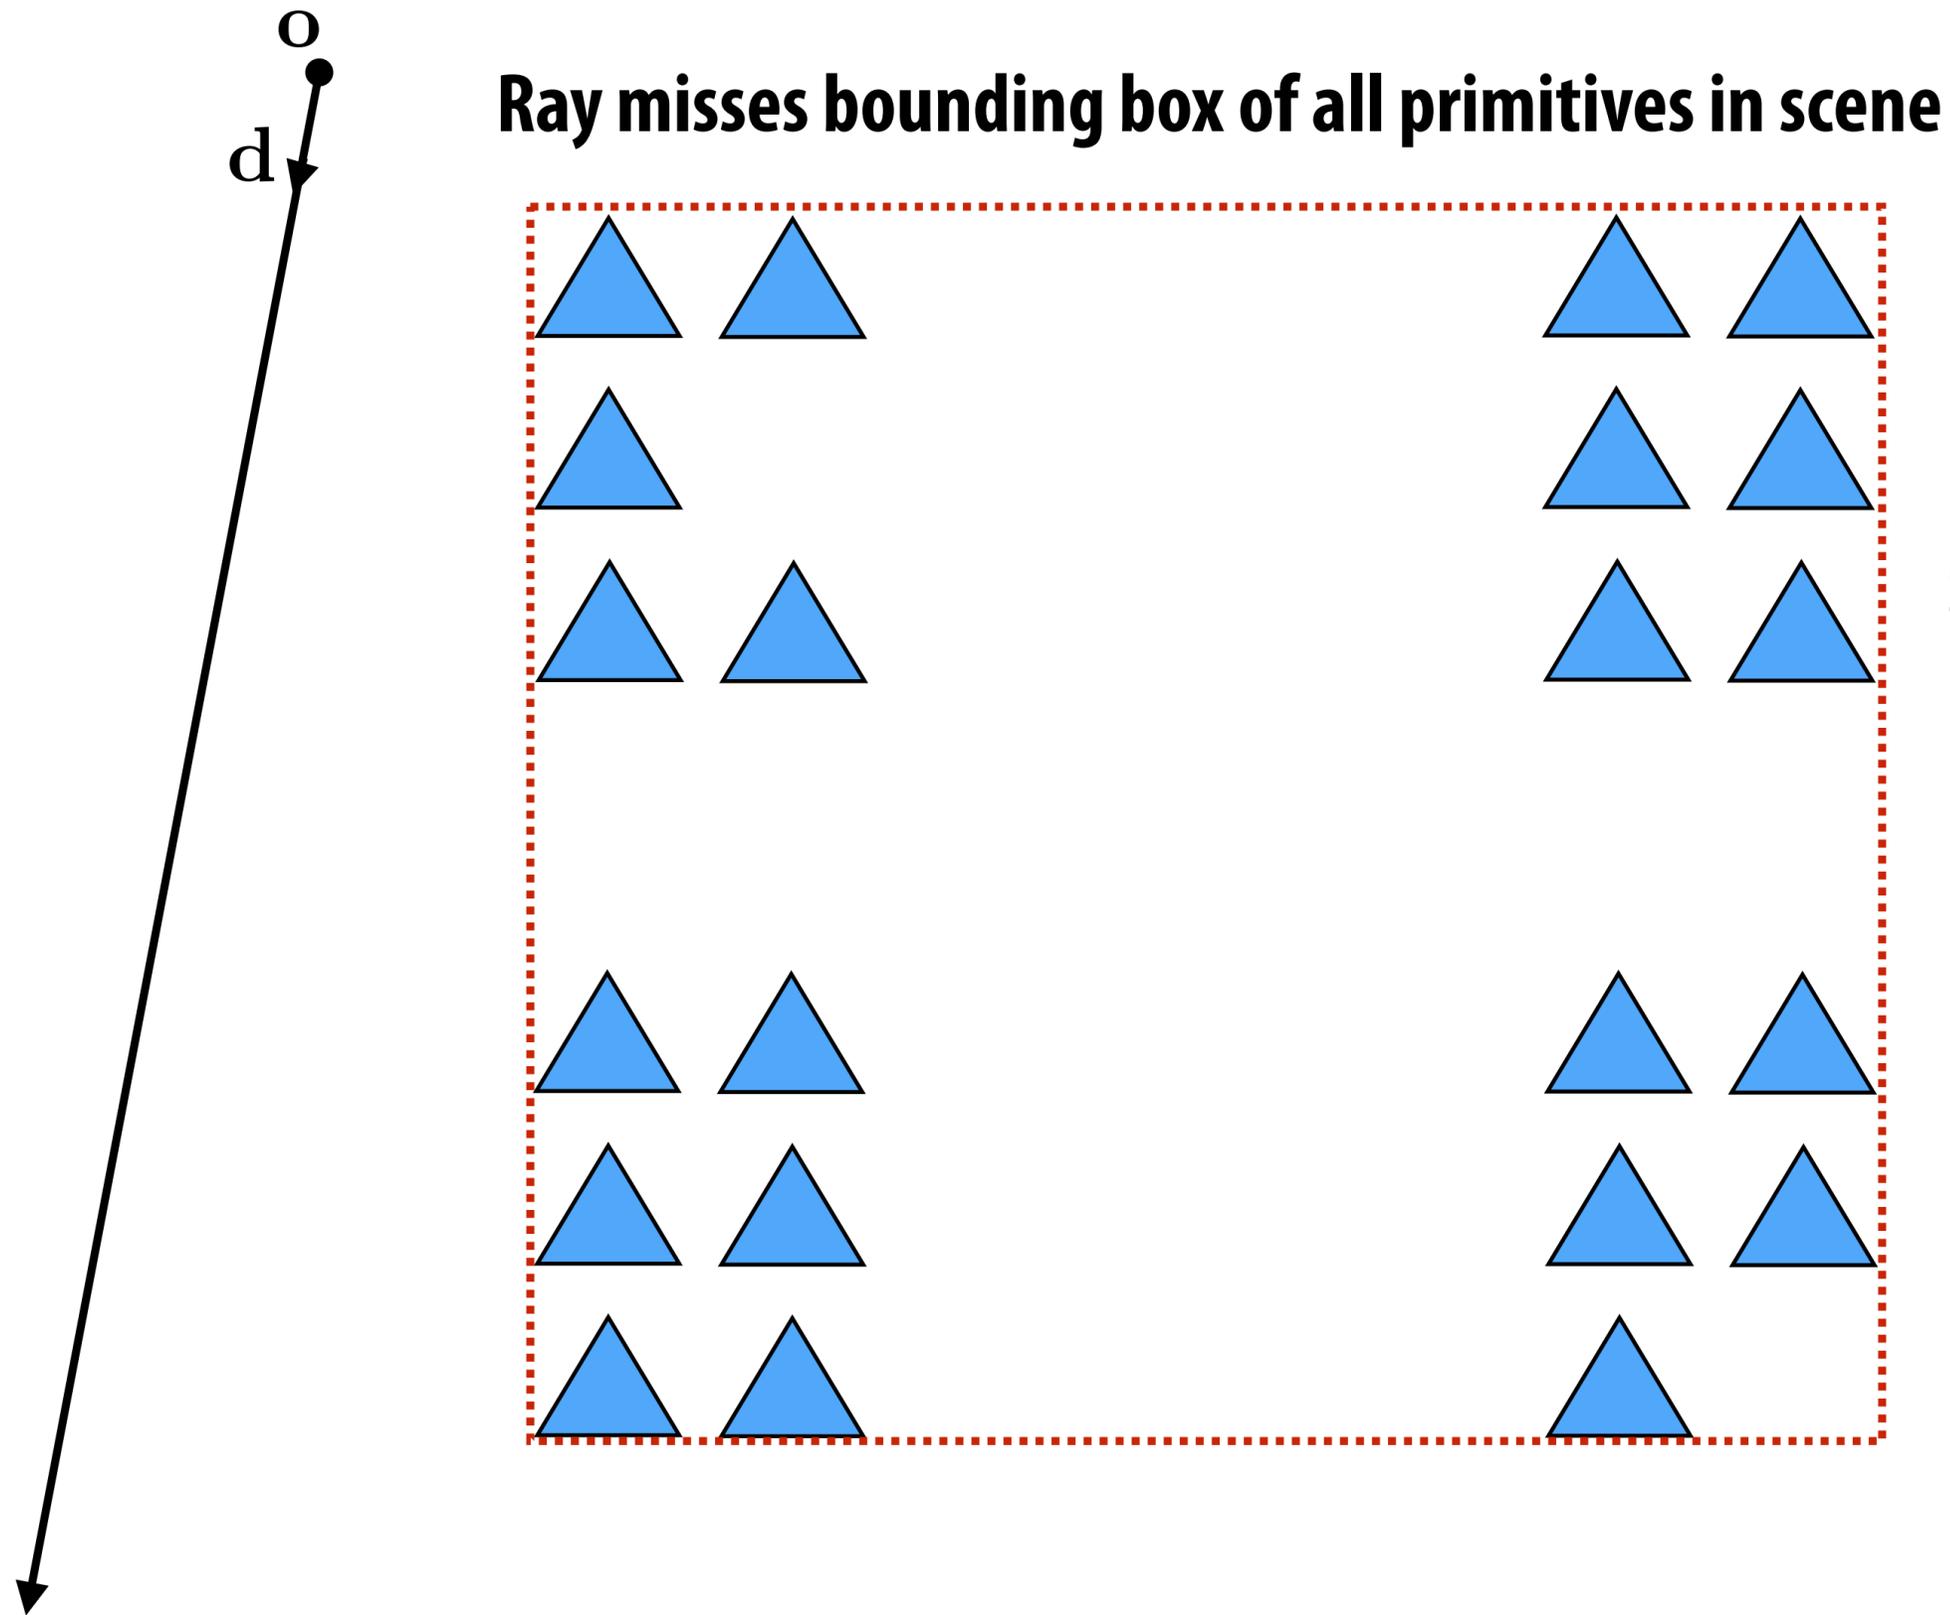
worse than before! :-)

...*much* better!

# Can we also reorganize scene primitives to enable fast ray-scene intersection queries?



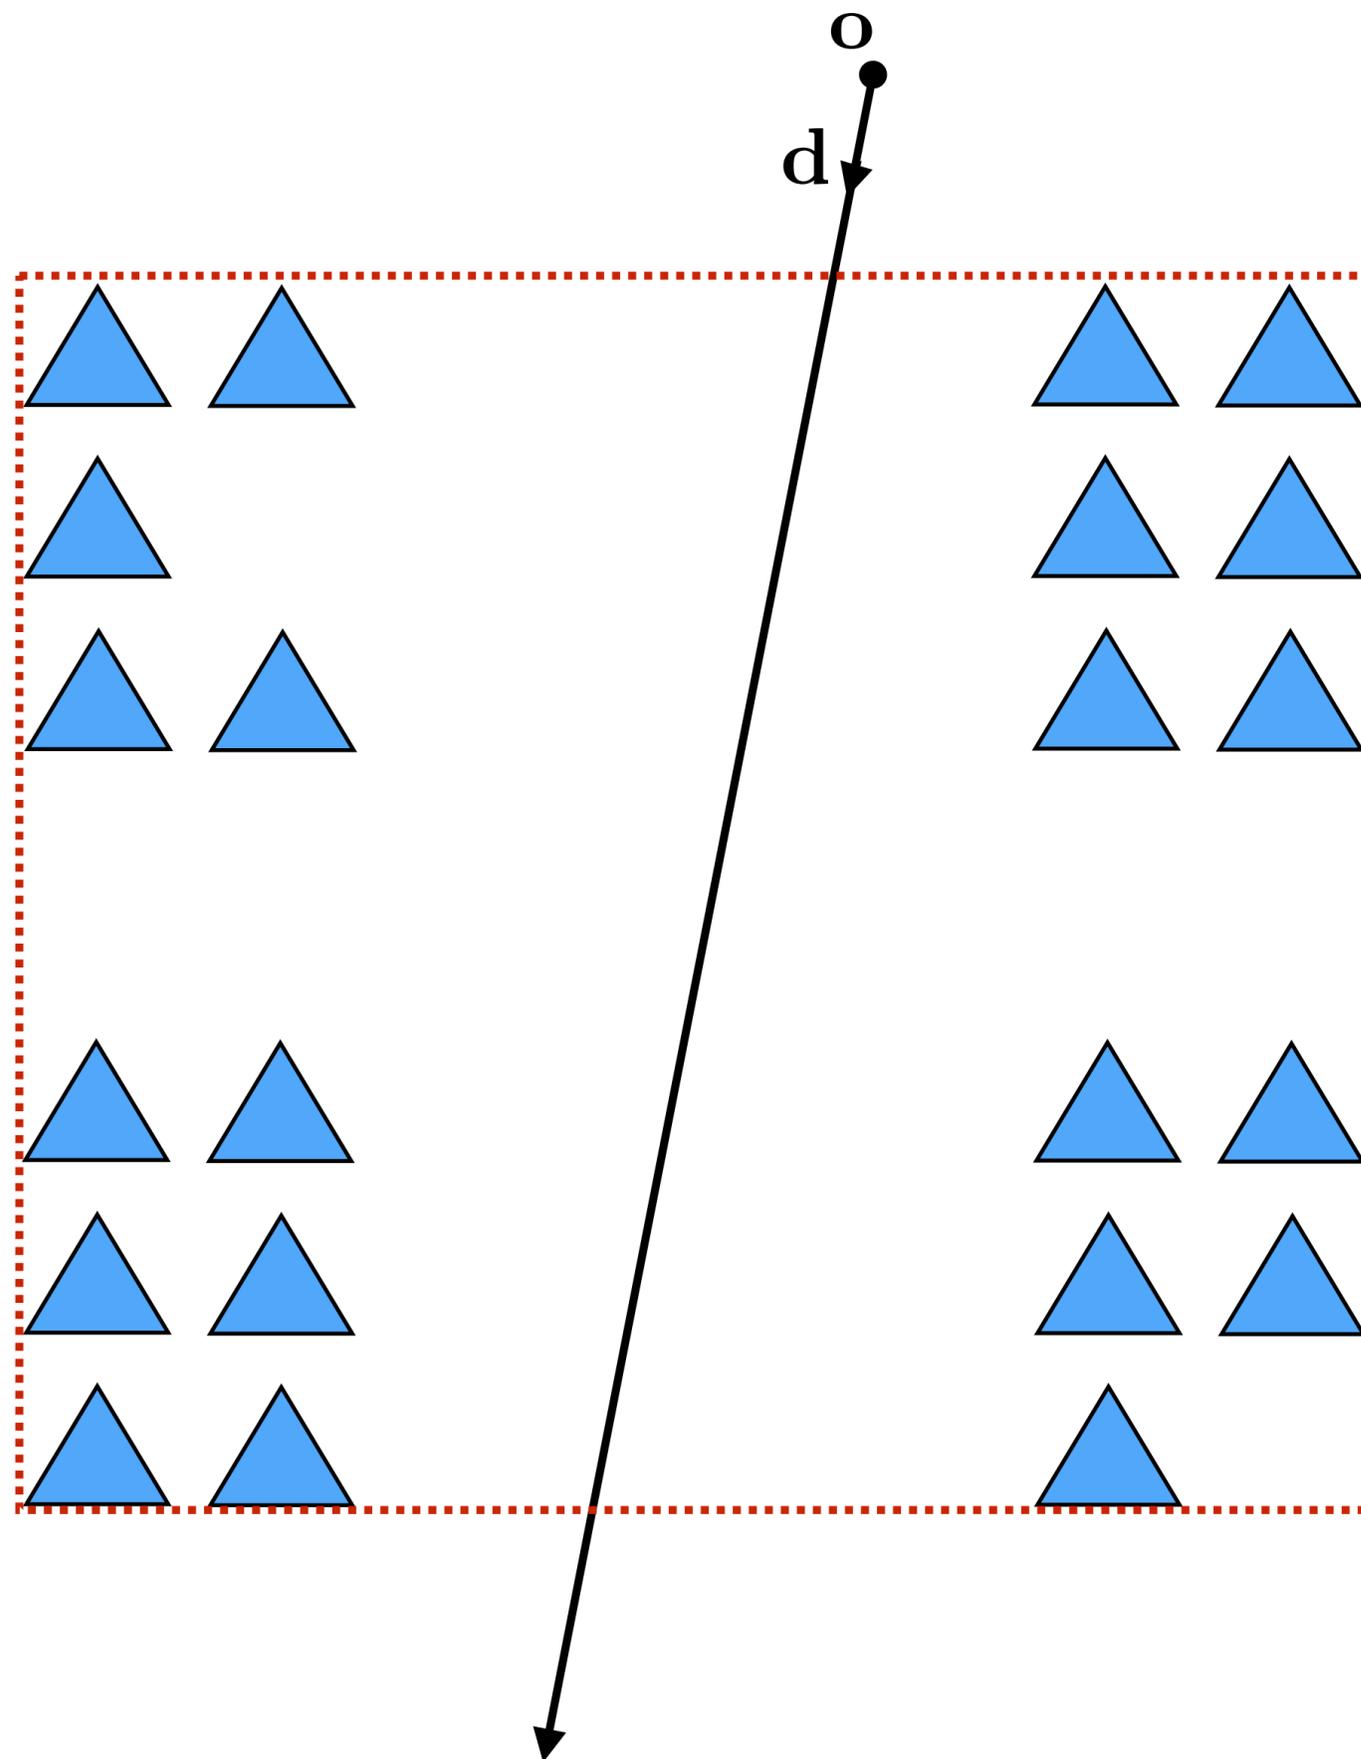
# Simple case



**Cost (misses box):**  
**preprocessing:  $O(n)$**   
**ray-box test:  $O(1)$**   
**amortized cost\*:  $O(1)$**

**\*over *many* ray-scene intersection tests**

# Another (should be) simple case



**Cost (hits box):**

**preprocessing:  $O(n)$**

**ray-box test:  $O(1)$**

**triangle tests:  $O(n)$**

**amortized cost\*:  $O(n)$**

**Still no better than  
naïve algorithm  
(test all triangles)!**

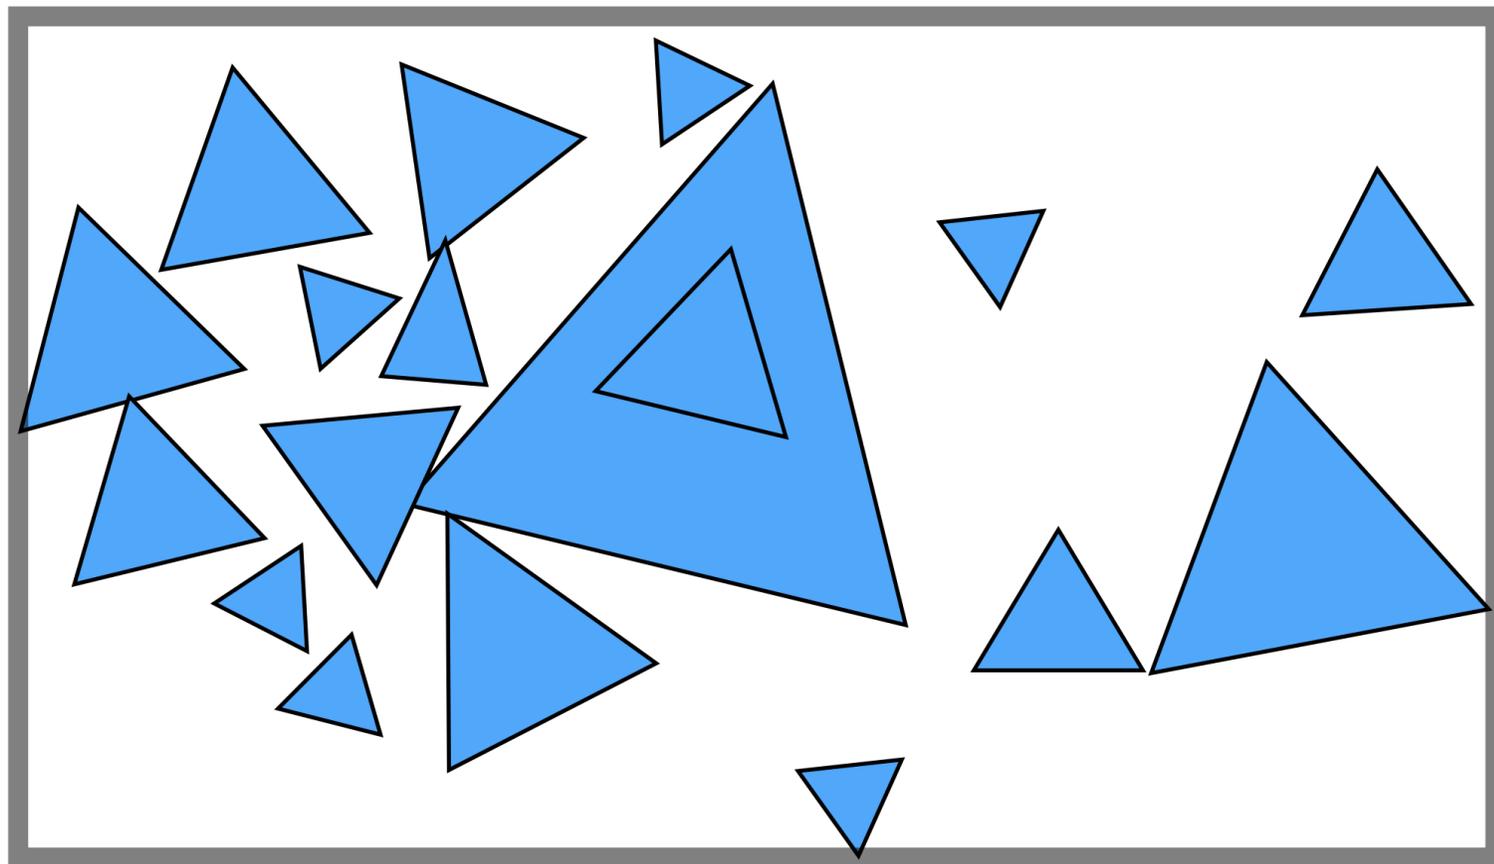
**\*over *many* ray-scene intersection tests**

**Q: How can we do better?**

**A: Apply this strategy hierarchically**

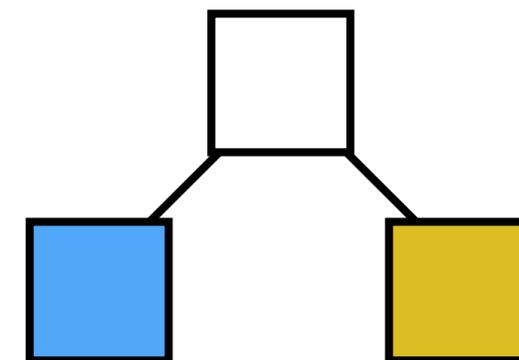
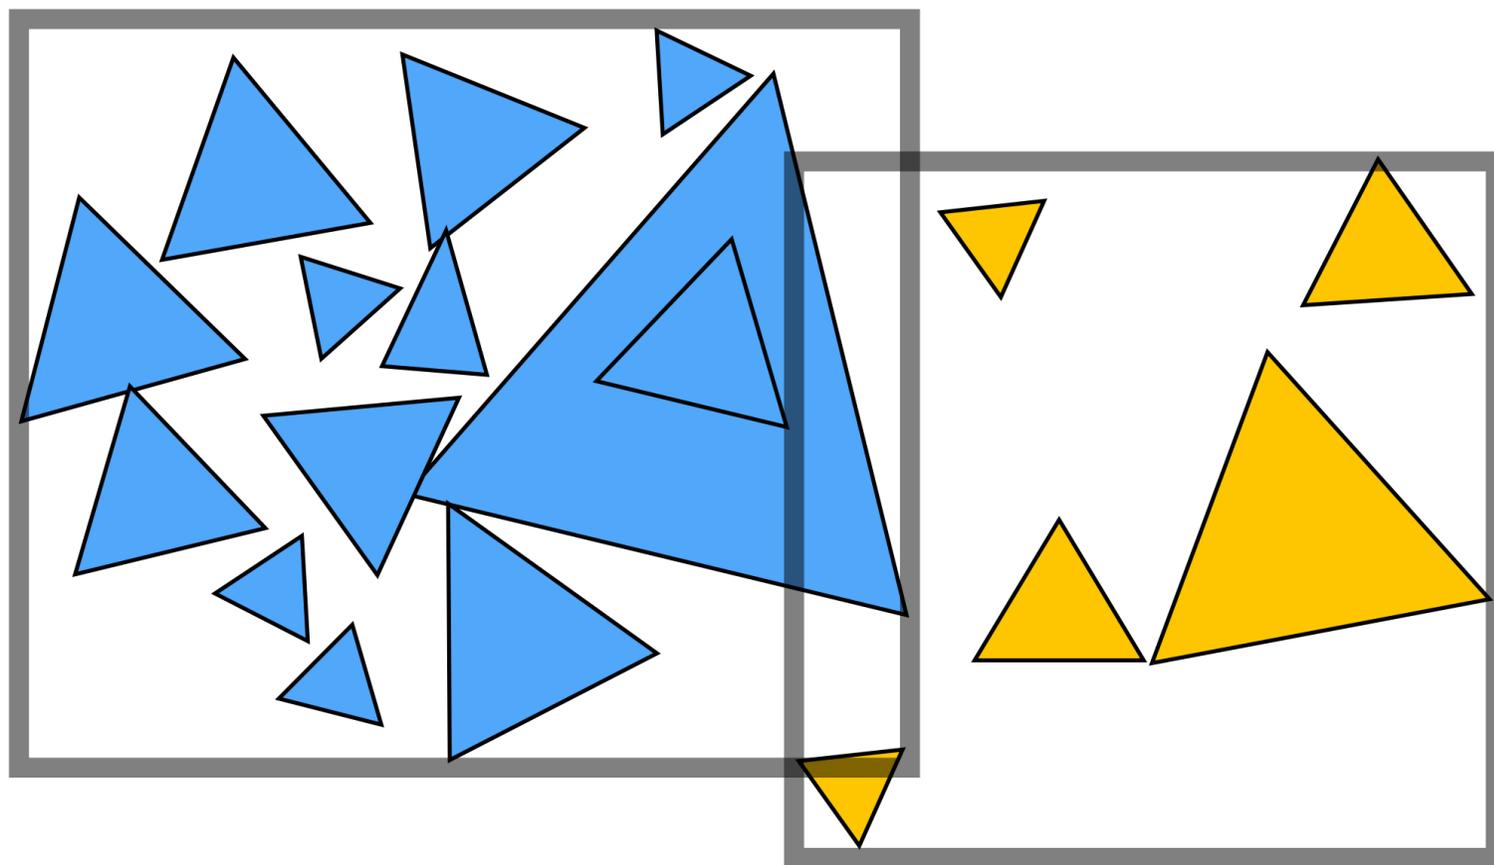
# Bounding volume hierarchy (BVH)

Root → 

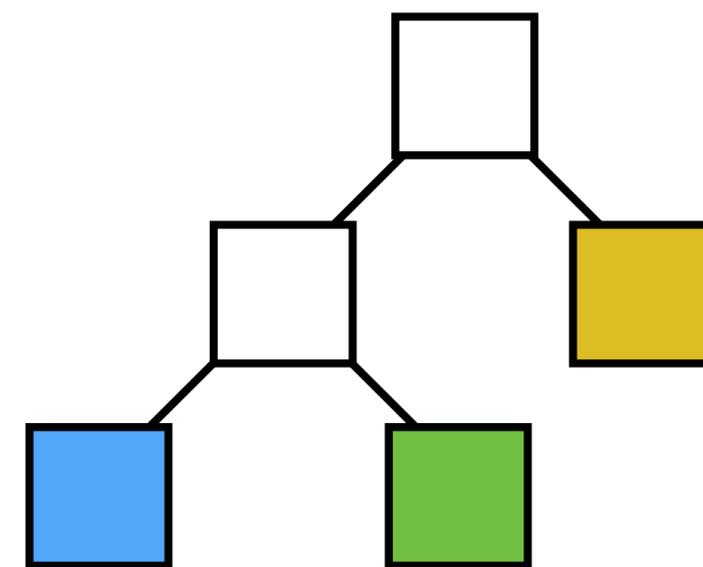
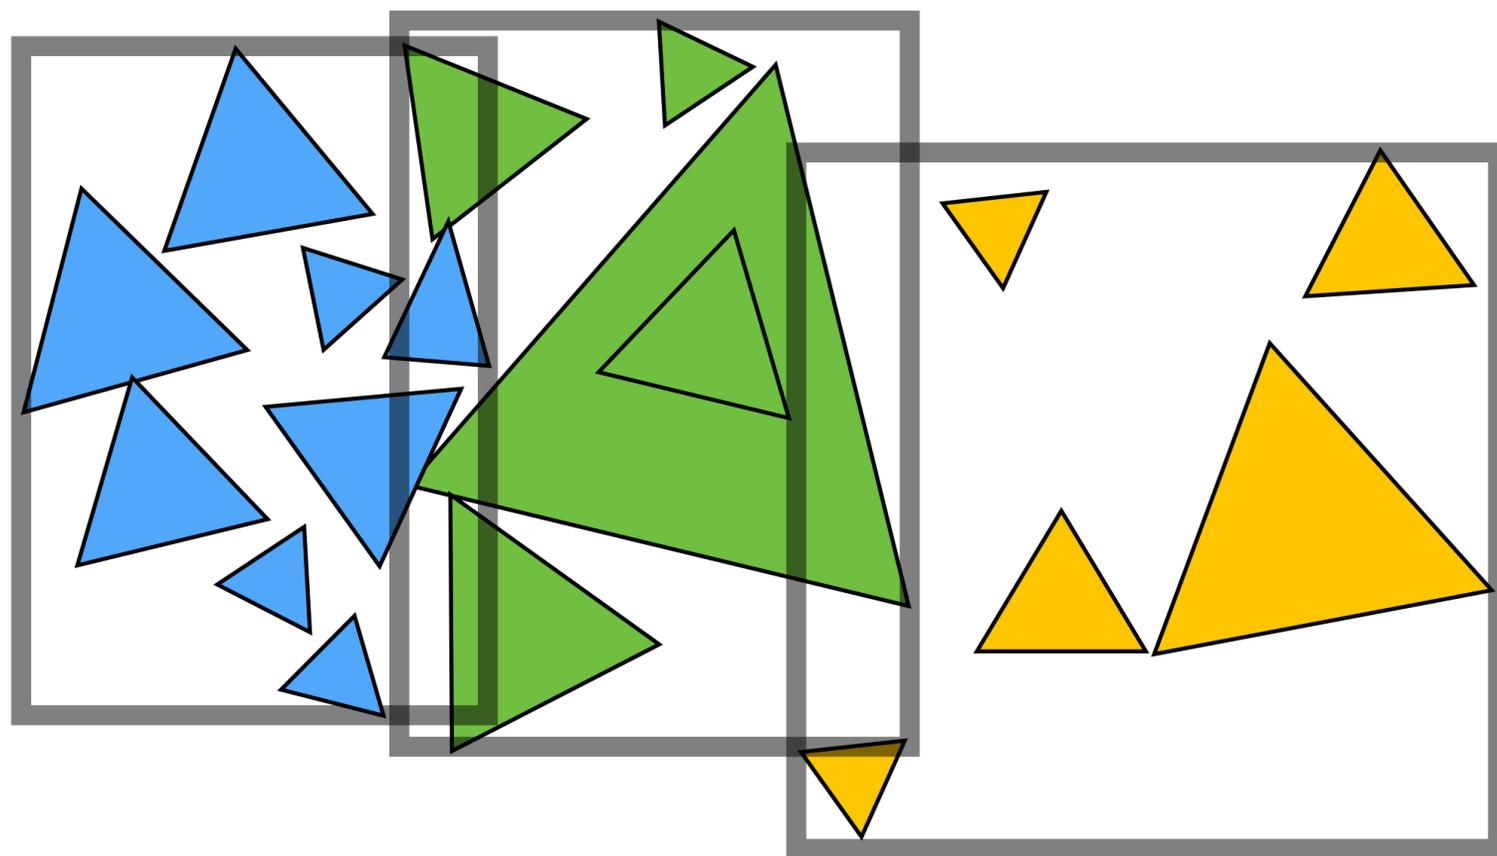


# Bounding volume hierarchy (BVH)

- BVH partitions each node's primitives into disjoint sets
  - Note: the sets can overlap in space (see example below)

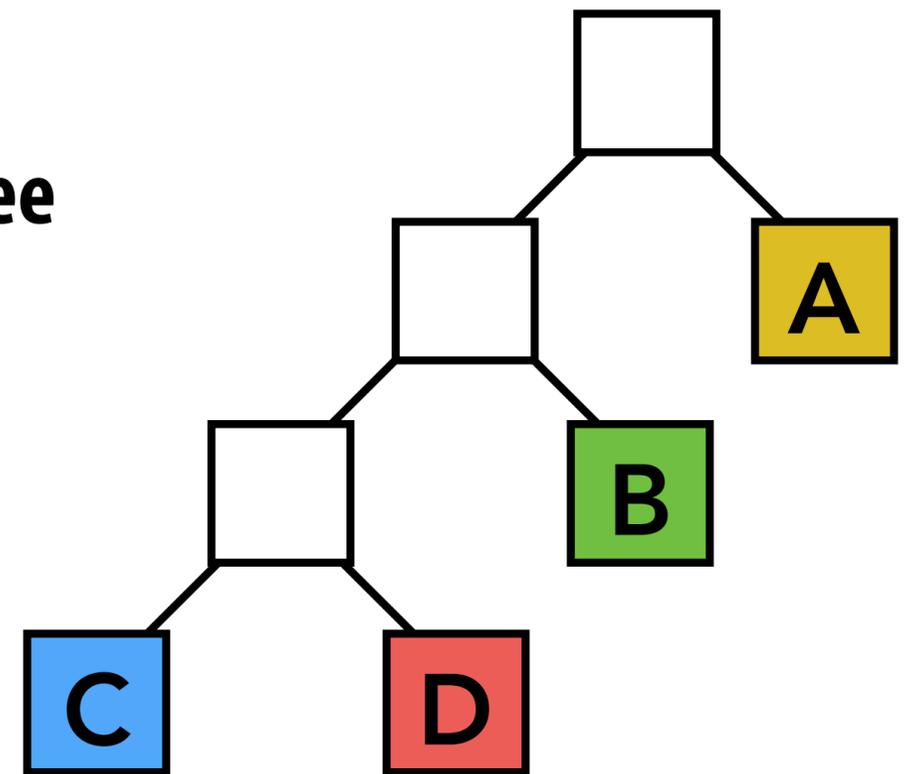
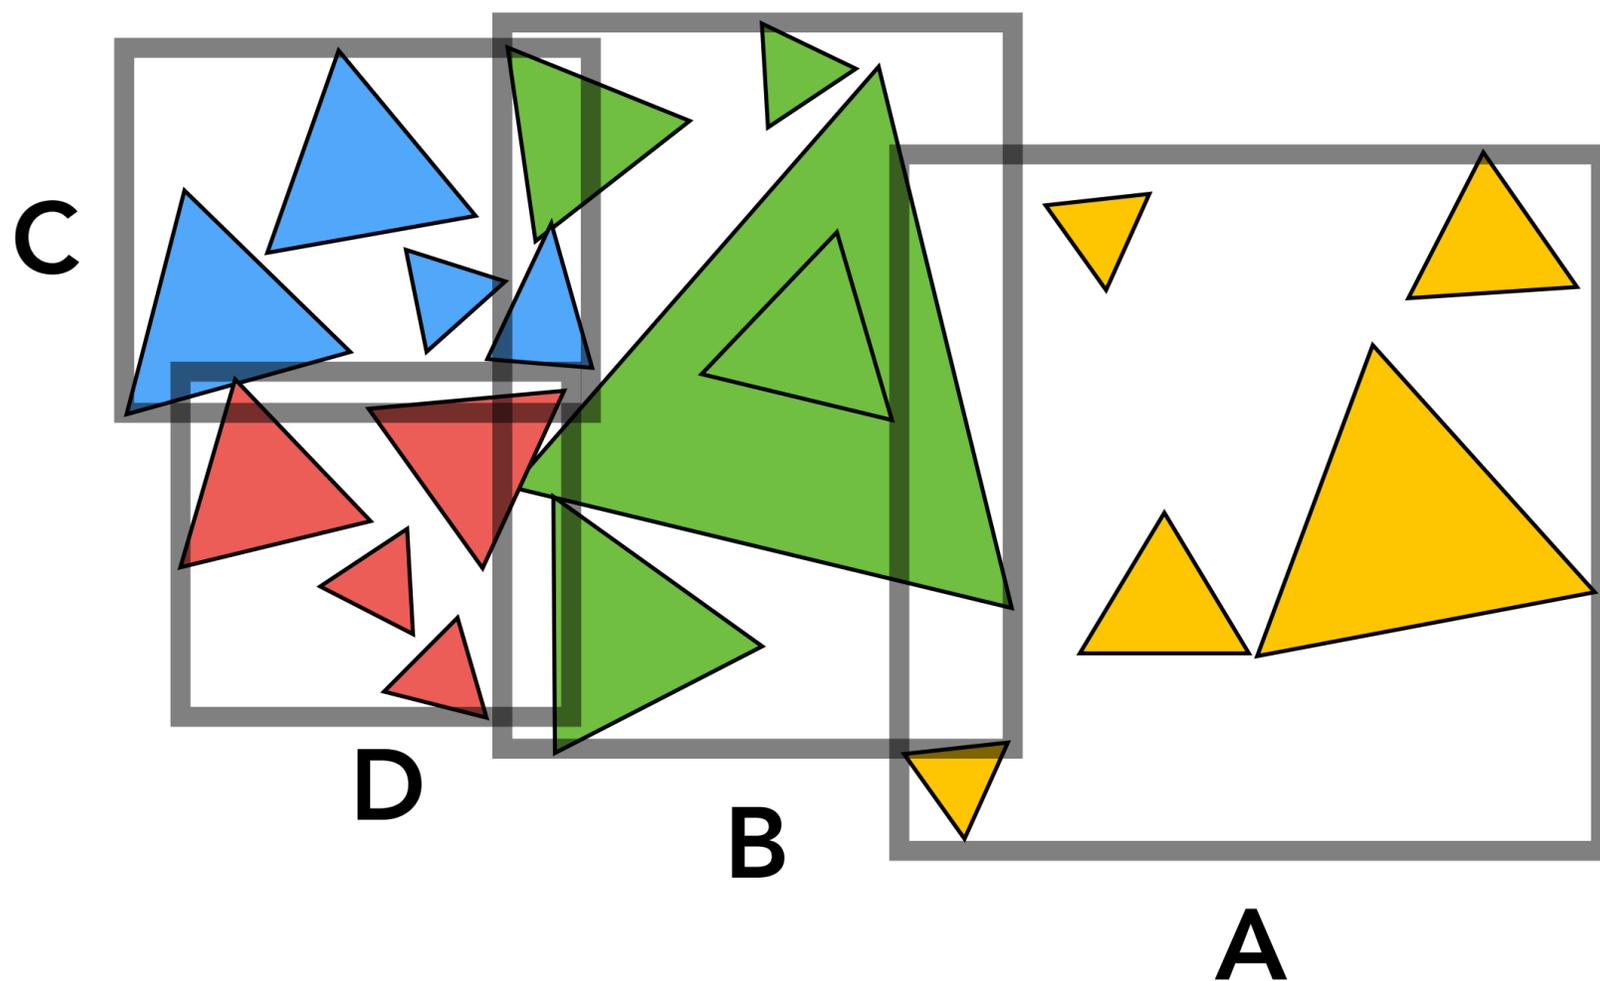


# Bounding volume hierarchy (BVH)

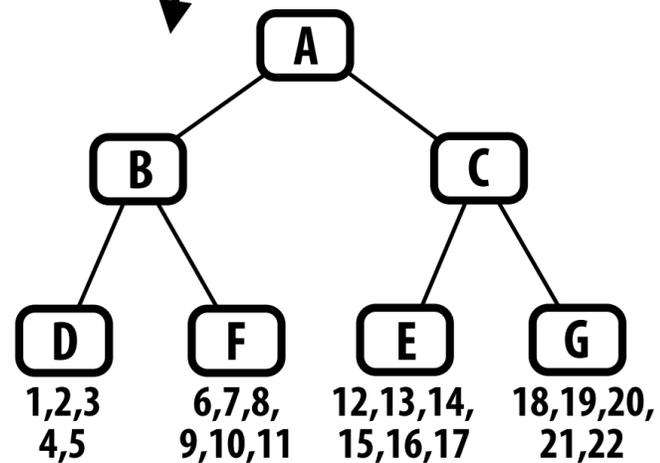
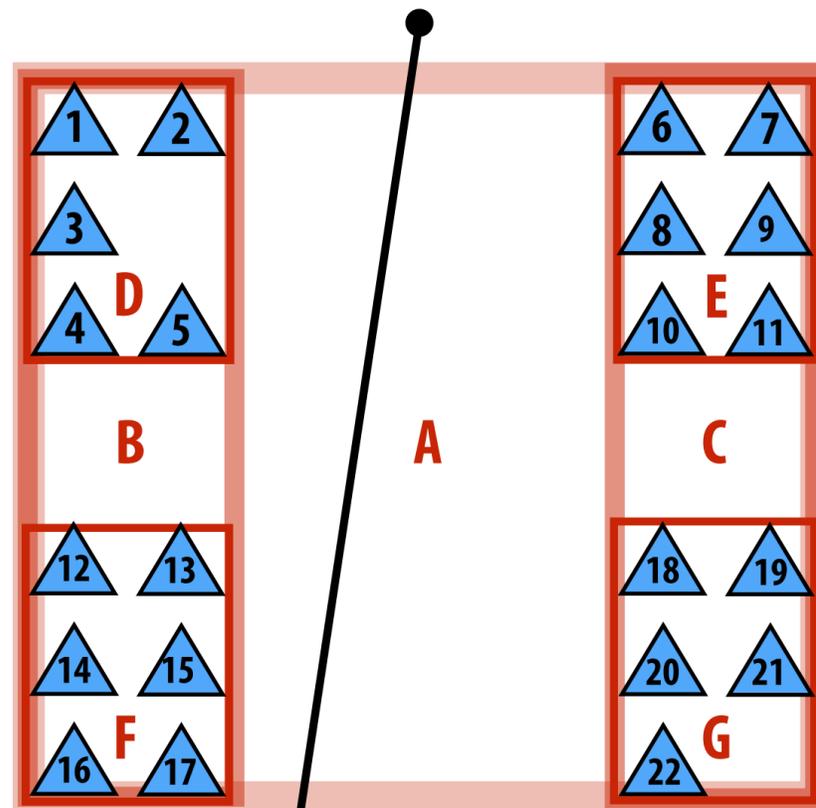
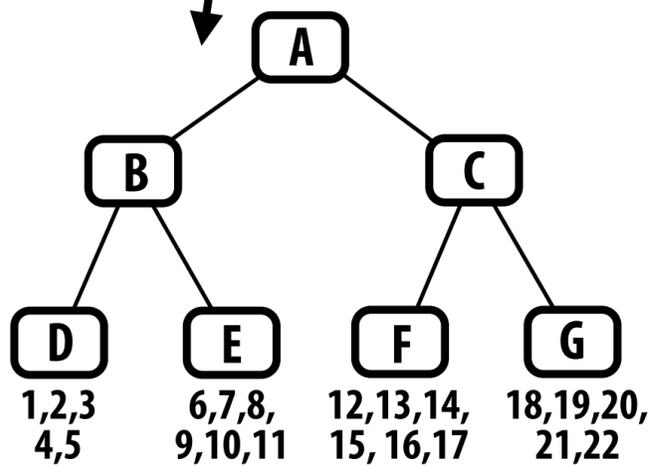
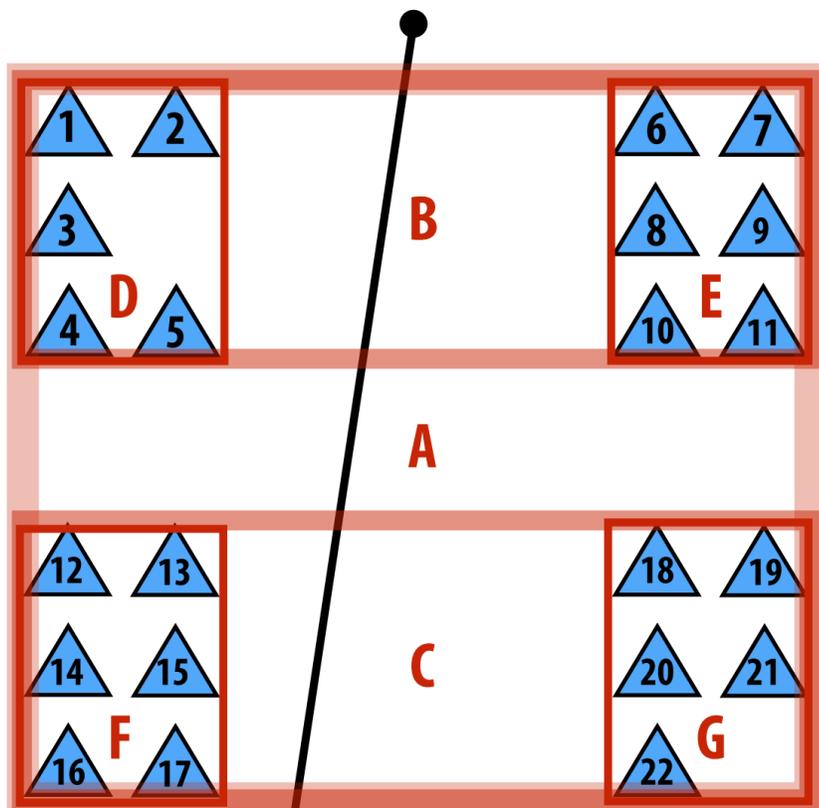


# Bounding volume hierarchy (BVH)

- Leaf nodes:
  - Contain *small* list of primitives
- Interior nodes:
  - Proxy for a *large* subset of primitives
  - Stores bounding box for all primitives in subtree



# Bounding volume hierarchy (BVH)



Left: two different BVH organizations of the same scene containing 22 primitives.

Is one BVH better than the other?

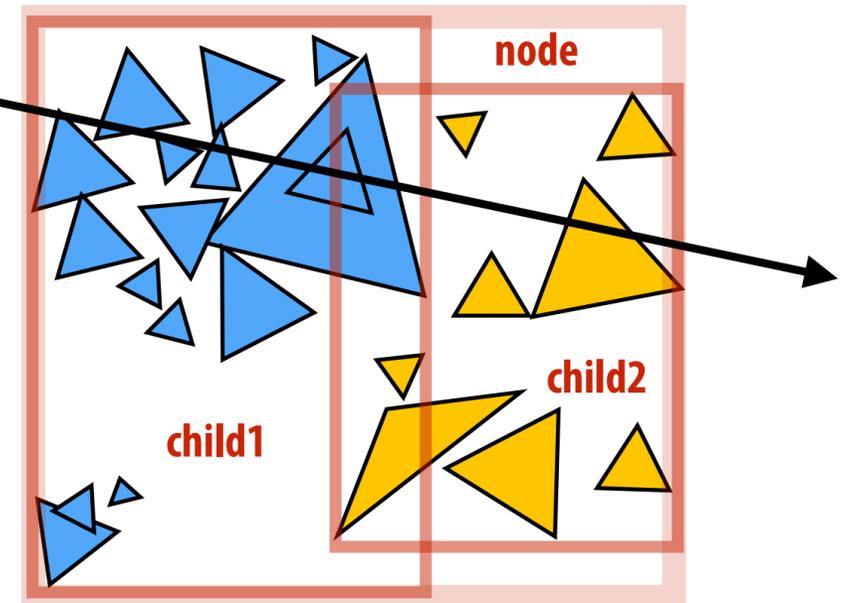
# Ray-scene intersection using a BVH

```
struct BVHNode {
    bool leaf; // true if node is a leaf
    BBox bbox; // min/max coords of enclosed primitives
    BVHNode* child1; // "left" child (could be NULL)
    BVHNode* child2; // "right" child (could be NULL)
    Primitive* primList; // for leaves, stores primitives
};
```

```
struct HitInfo {
    Primitive* prim; // which primitive did the ray hit?
    float t; // at what t value along ray?
};
```

```
void find_closest_hit(Ray* ray, BVHNode* node, HitInfo* closest) {
    HitInfo hit = intersect(ray, node->bbox); // test ray against node's bounding box
    if (hit.prim == NULL || hit.t > closest.t)
        return; // don't update the hit record

    if (node->leaf) {
        for (each primitive p in node->primList) {
            hit = intersect(ray, p);
            if (hit.prim != NULL && hit.t < closest.t) {
                closest.prim = p;
                closest.t = t;
            }
        }
    }
    else {
        find_closest_hit(ray, node->child1, closest);
        find_closest_hit(ray, node->child2, closest);
    }
}
```

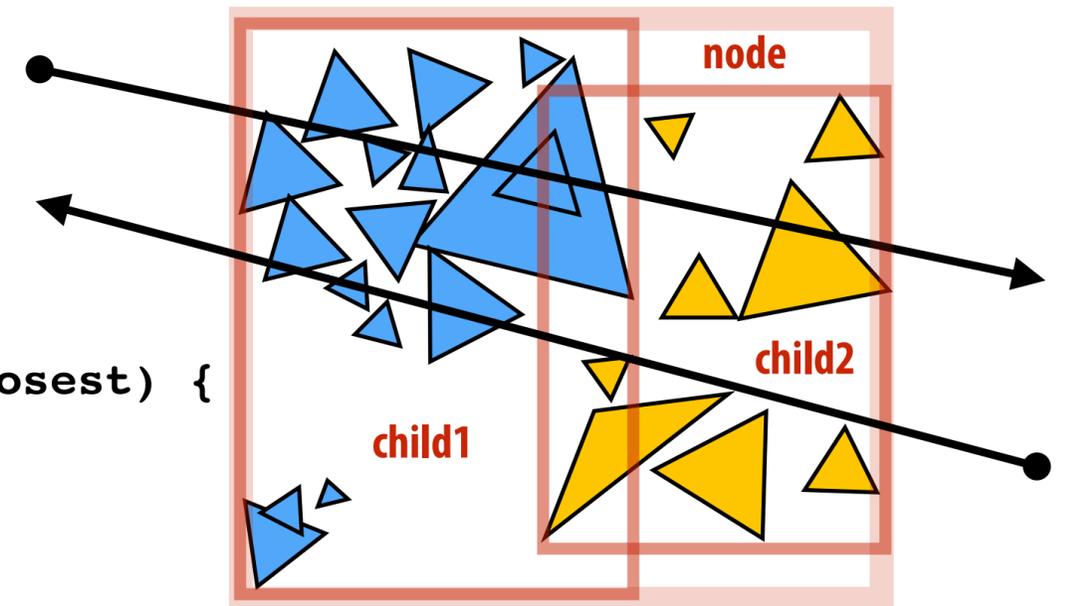


**How could this occur?**

# Improvement: “front-to-back” traversal

**Invariant: only call `find_closest_hit()` if ray intersects bbox of node.**

```
void find_closest_hit(Ray* ray, BVHNode* node, HitInfo* closest) {  
  
    if (node->leaf) {  
        for (each primitive p in node->primList) {  
            hit = intersect(ray, p);  
            if (hit.prim != NULL && t < closest.t) {  
                closest.prim = p;  
                closest.t = t;  
            }  
        }  
    }  
    else {  
        HitInfo hit1 = intersect(ray, node->child1->bbox);  
        HitInfo hit2 = intersect(ray, node->child2->bbox);  
  
        NVHNode* first = (hit1.t <= hit2.t) ? child1 : child2;  
        NVHNode* second = (hit1.t <= hit2.t) ? child2 : child1;  
  
        find_closest_hit(ray, first, closest);  
        if (second child's t is closer than closest.t)  
            find_closest_hit(ray, second, closest); // why might we still need to do this?  
    }  
}
```

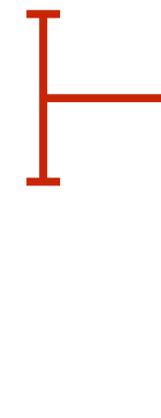


**“Front to back” traversal. Traverse to closest child node first. Why?**

# Aside: another type of query: any hit

Sometimes it's useful to know if the ray hits ANY primitive in the scene at all (don't care about distance to first hit)

```
bool find_any_hit(Ray* ray, BVHNode* node) {  
  
    if (!intersect(ray, node->bbox))  
        return false;  
  
    if (node->leaf) {  
        for (each primitive p in node->primList) {  
            hit = intersect(ray, p);  
            if (hit.prim)  
                return true;  
        }  
    } else {  
        return ( find_closest_hit(ray, node->child1, closest) ||  
                find_closest_hit(ray, node->child2, closest) );  
    }  
}
```



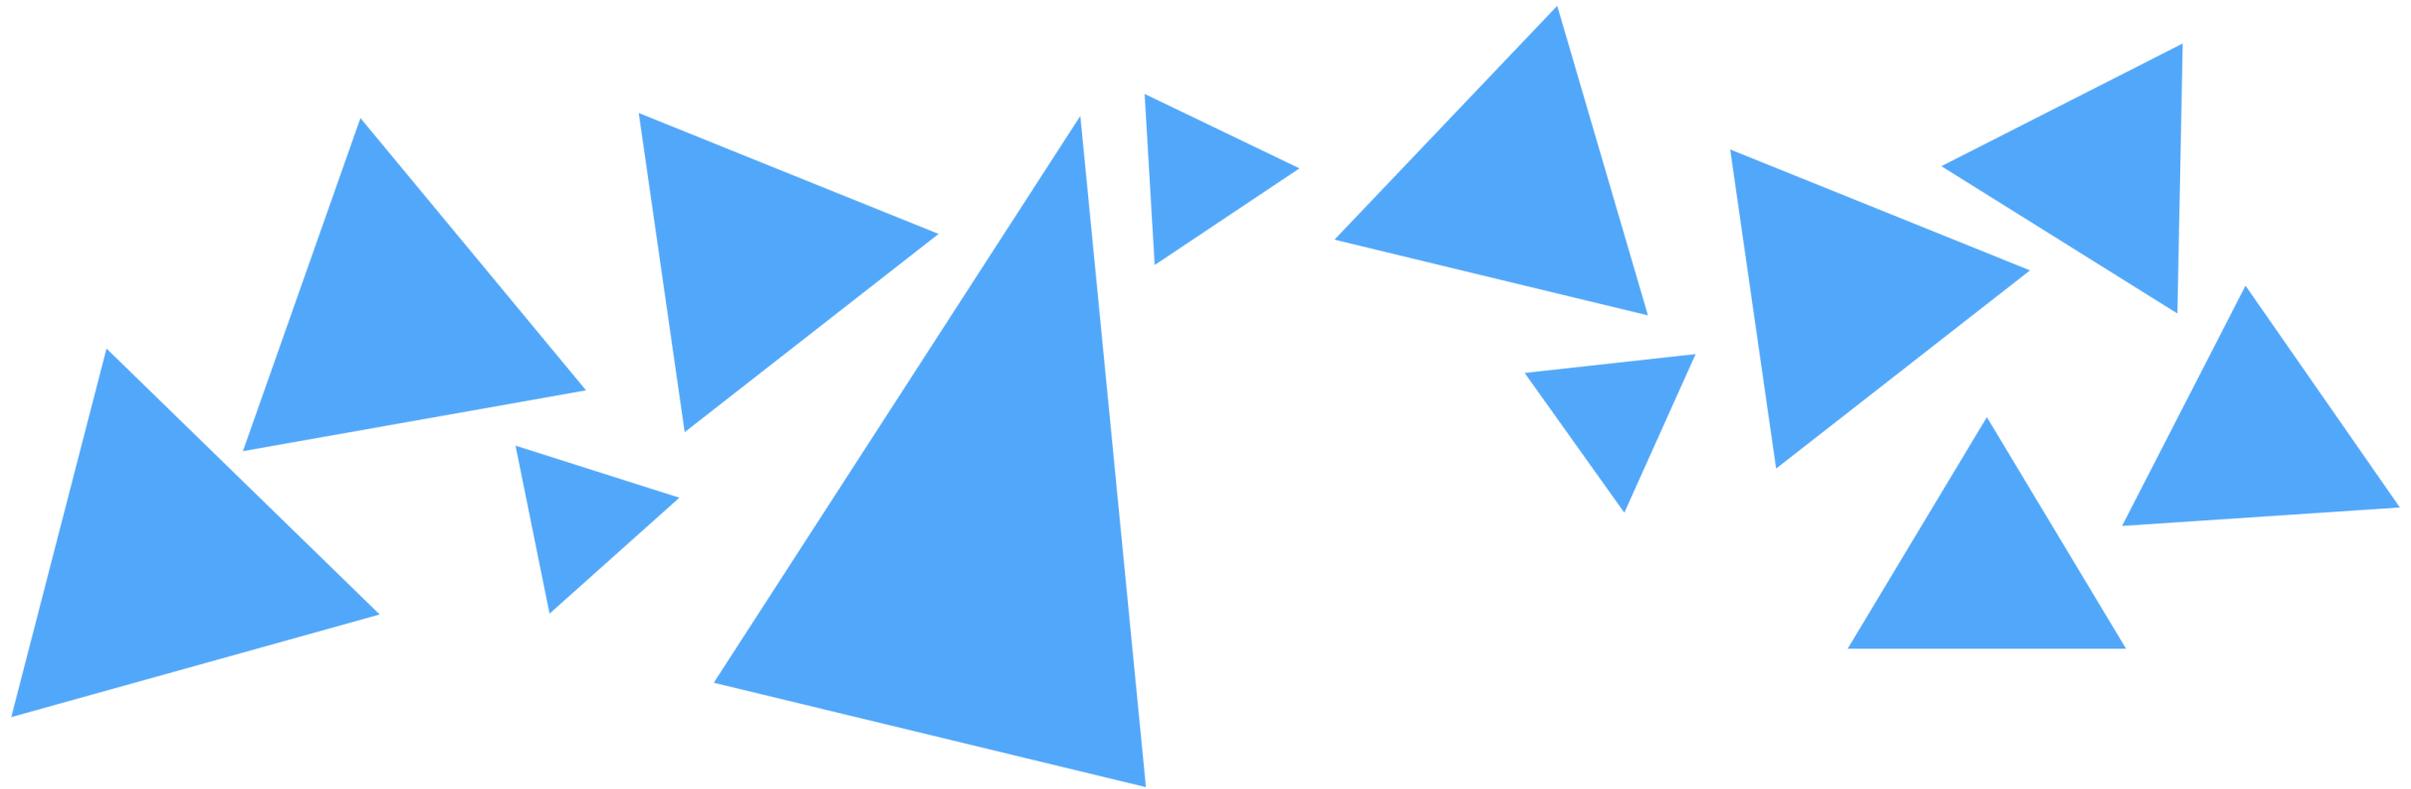
Interesting question of which child to enter first. How might you make a good decision?

**For a given set of primitives, there are  
many possible BVHs**

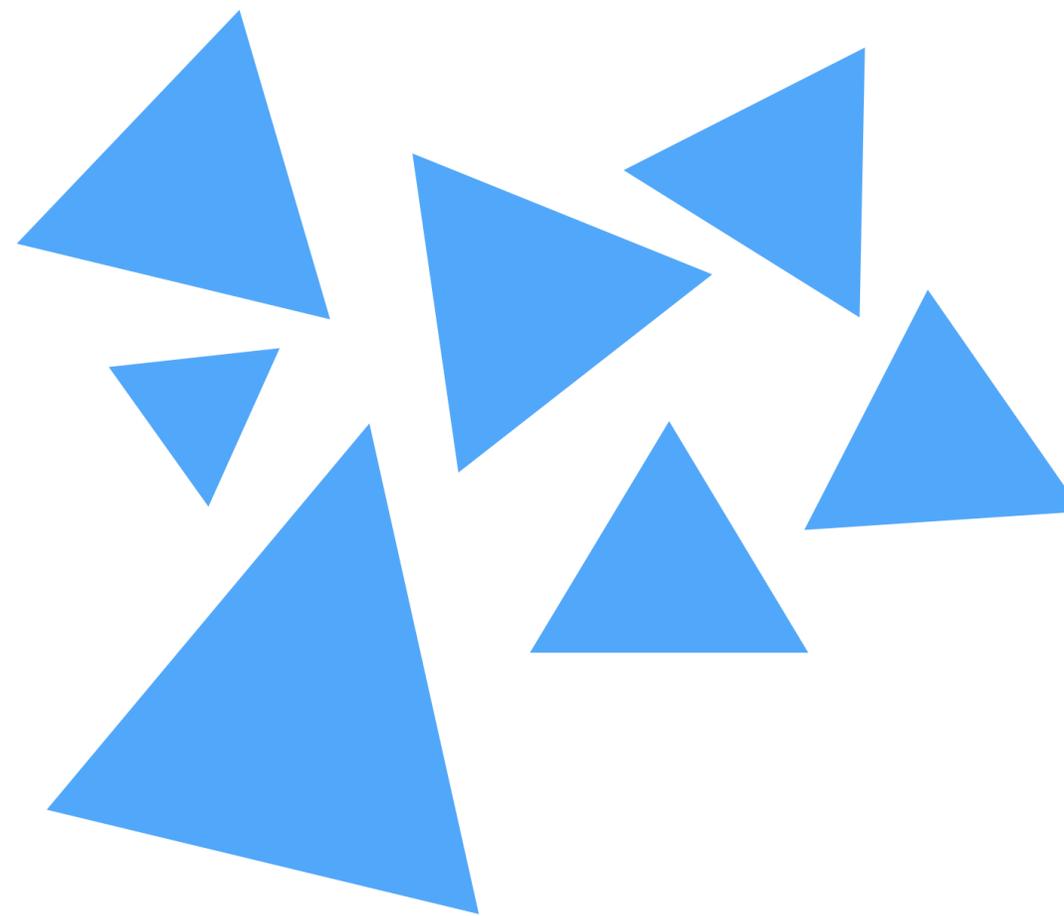
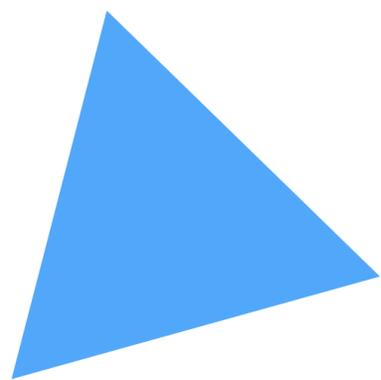
**( $2^{N/2}$  ways to partition  $N$  primitives into two groups)**

**Q: How do we build a high-quality BVH?**

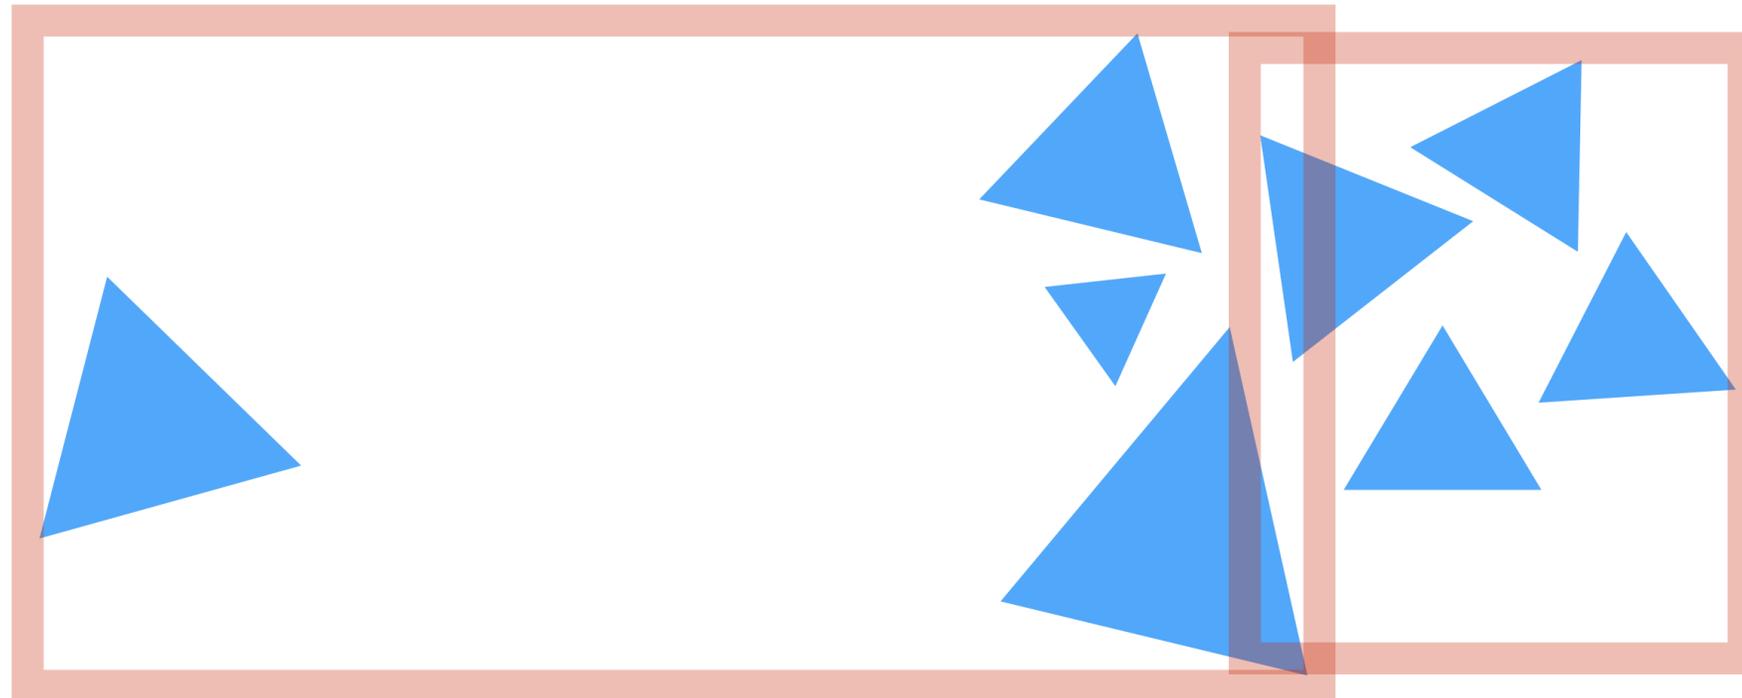
# How would you partition these triangles into two groups?



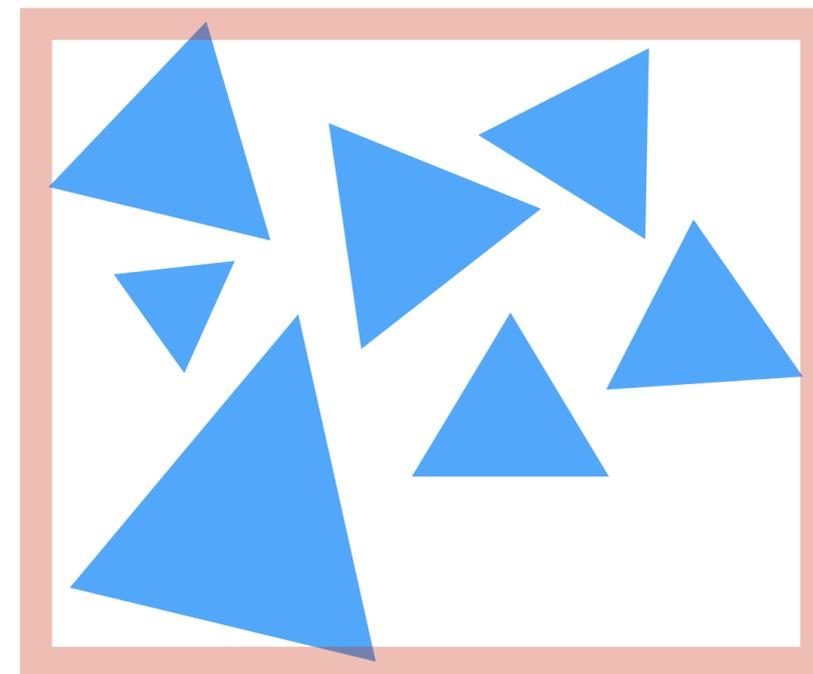
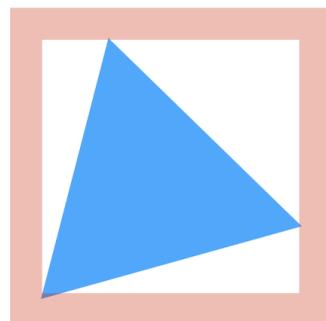
# What about these?



# Intuition about a “good” partition?



**Partition into child nodes with equal numbers of primitives**



**Better partition**

**Intuition: want small bounding boxes (minimize overlap between children, avoid bboxes with empty space)**

# What are we really trying to do?

A good partitioning minimizes the cost of finding the closest intersection of a ray with primitives in the node.

If a node is a leaf node (no partitioning):

$$C = \sum_{i=1}^N C_{\text{isect}}(i)$$
$$= N C_{\text{isect}}$$

Where  $C_{\text{isect}}(i)$  is the cost of ray-primitive intersection for primitive  $i$  in the node.

(Common to assume all primitives have the same cost)

# Cost of making a partition

The expected cost of ray-node intersection, given that the node's primitives are partitioned into child sets A and B is:

$$C = C_{\text{trav}} + p_A C_A + p_B C_B$$

$C_{\text{trav}}$  is the cost of traversing an interior node (e.g., load data, bbox intersection check)

$C_A$  and  $C_B$  are the costs of intersection with the resultant child subtrees

$p_A$  and  $p_B$  are the probability a ray intersects the bbox of the child nodes A and B

**Primitive count is common approximation for child node costs:**

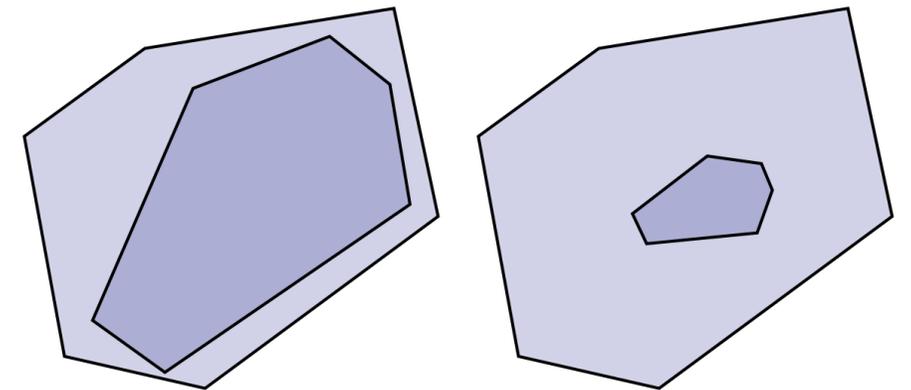
$$C = C_{\text{trav}} + p_A N_A C_{\text{isect}} + p_B N_B C_{\text{isect}}$$

**Remaining question: how do we get the probabilities  $p_A$ ,  $p_B$ ?**

# Estimating probabilities

- For convex object A inside convex object B, the probability that a random ray that hits B also hits A is given by the ratio of the surface areas  $S_A$  and  $S_B$  of these objects.

$$P(\text{hit } A | \text{hit } B) = \frac{S_A}{S_B}$$



Leads to surface area heuristic (SAH):

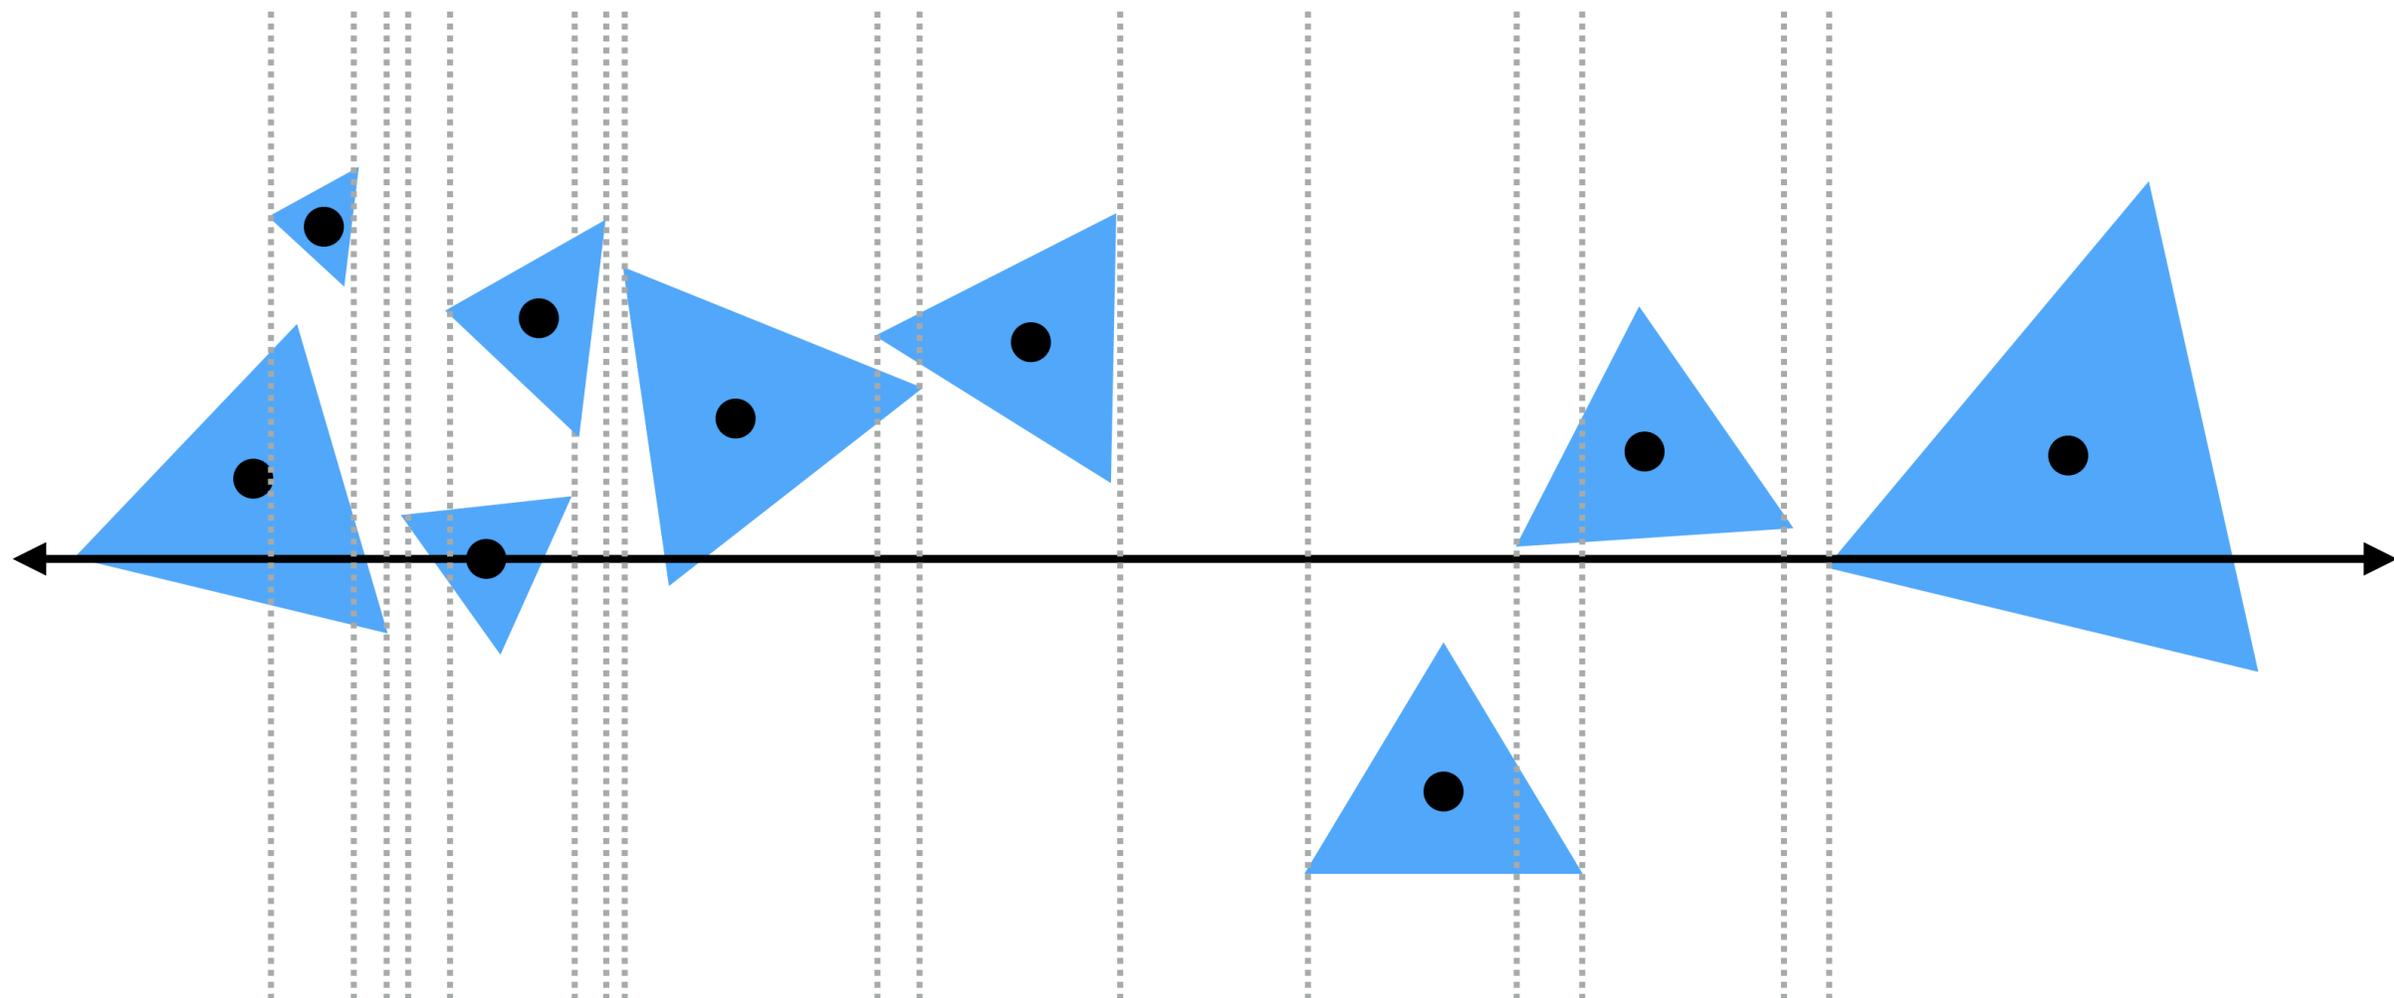
$$C = C_{\text{trav}} + \frac{S_A}{S_N} N_A C_{\text{isect}} + \frac{S_B}{S_N} N_B C_{\text{isect}}$$

**Assumptions of the SAH (which may not hold in practice!):**

- Rays are randomly distributed
- Rays are not occluded

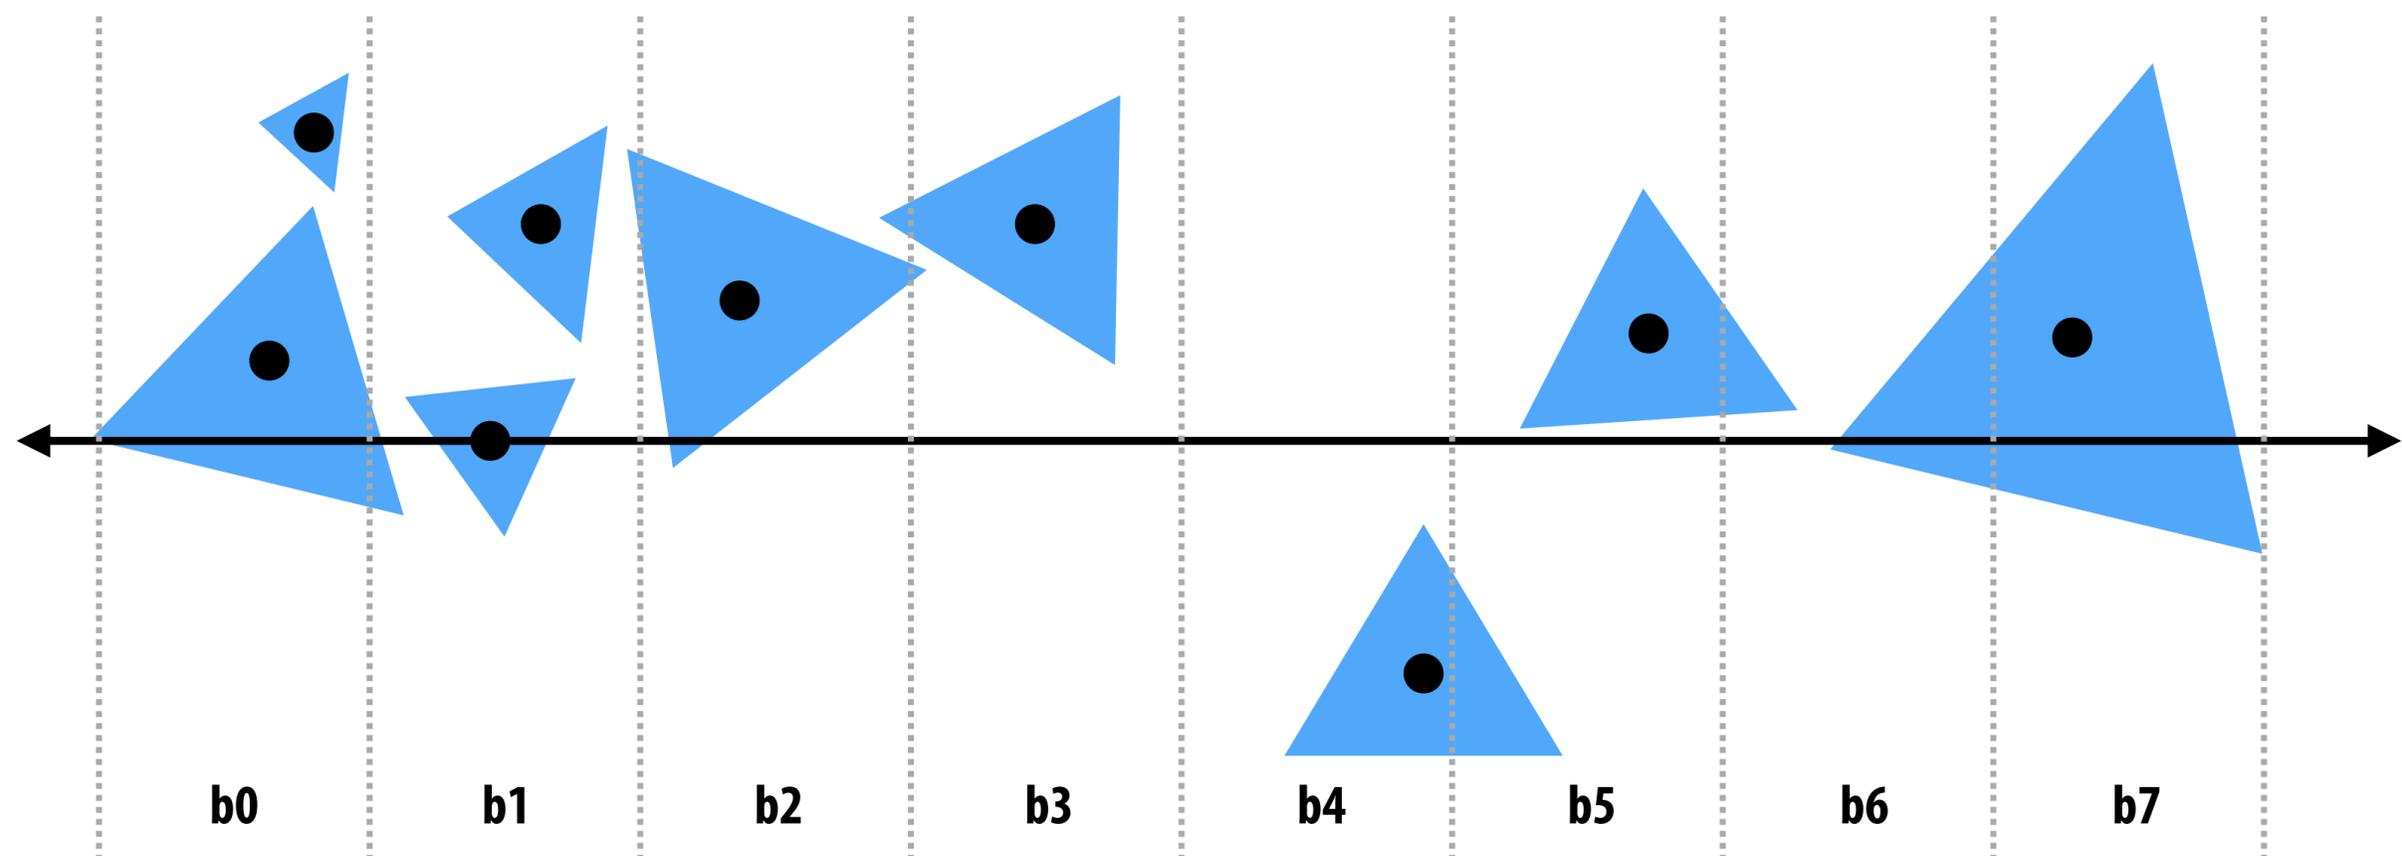
# Implementing partitions

- **Constrain search for good partitions to axis-aligned spatial partitions**
  - **Choose an axis; choose a split plane on that axis**
  - **Partition primitives by the side of splitting plane their centroid lies**
  - **SAH changes only when split plane moves past triangle boundary**
  - **Have to consider large number of possible split planes...  $O(\# \text{ objects})$**



# Efficiently implementing partitioning

- Efficient modern approximation: split spatial extent of primitives into  $B$  buckets ( $B$  is typically small:  $B < 32$ )



For each axis:  $x, y, z$ :  
initialize buckets

For each primitive  $p$  in node:

$b = \text{compute\_bucket}(p.\text{centroid})$

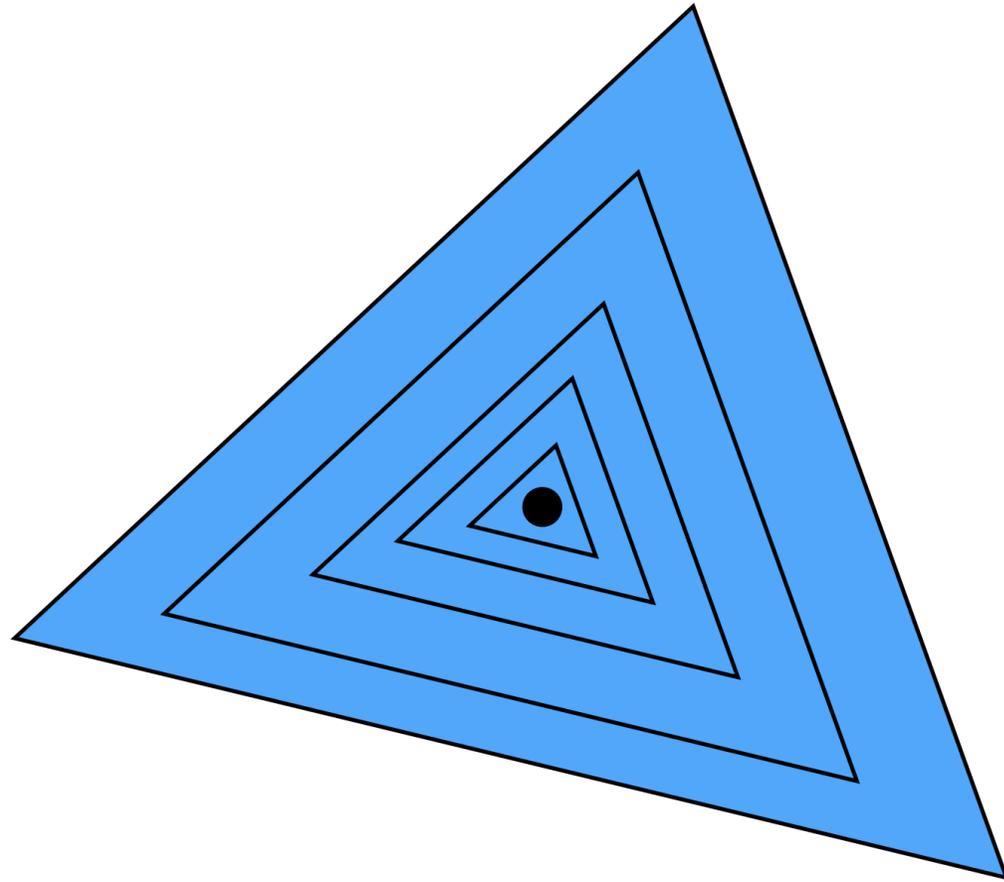
$b.\text{bbox.union}(p.\text{bbox});$

$b.\text{prim\_count}++;$

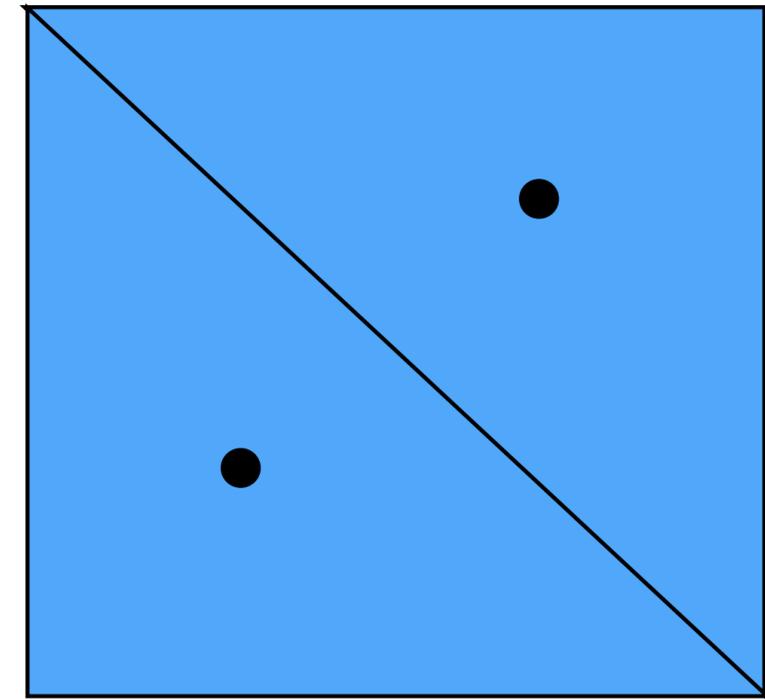
For each of the  $B-1$  possible partitioning planes evaluate SAH

Recurse on lowest cost partition found (or make node a leaf)

# Troublesome cases



**All primitives with same centroid (all primitives end up in same partition)**

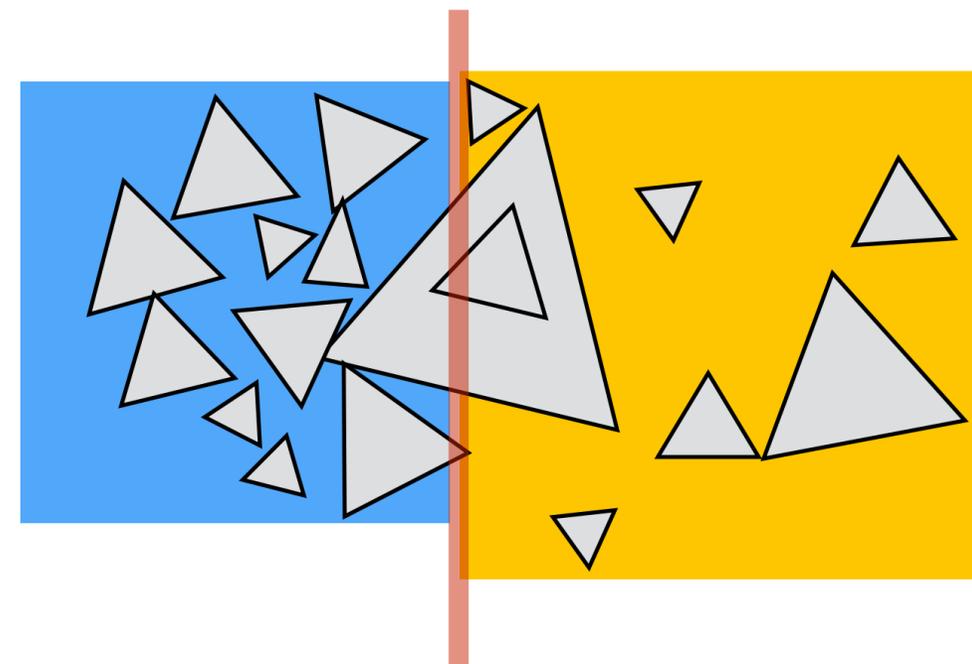
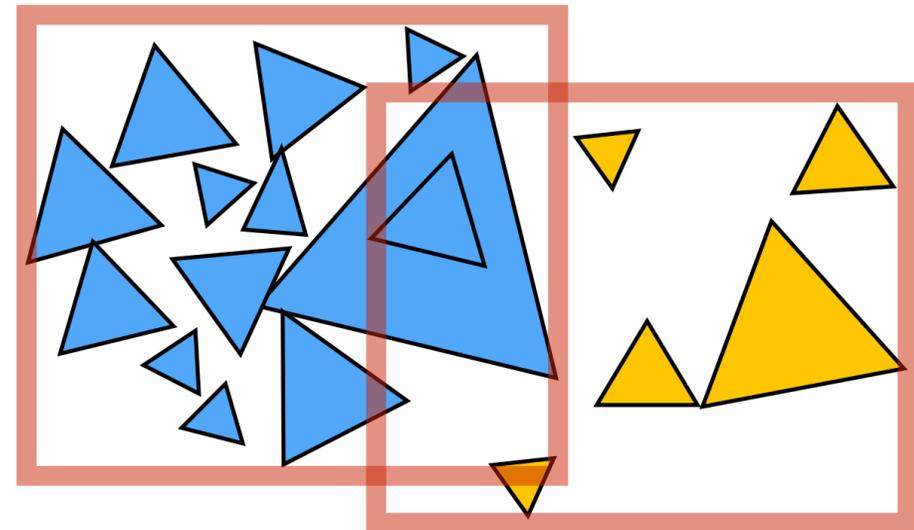


**All primitives with same bbox (ray often ends up visiting both partitions)**

**In general, different strategies may work better for different types of geometry / different distributions of primitives...**

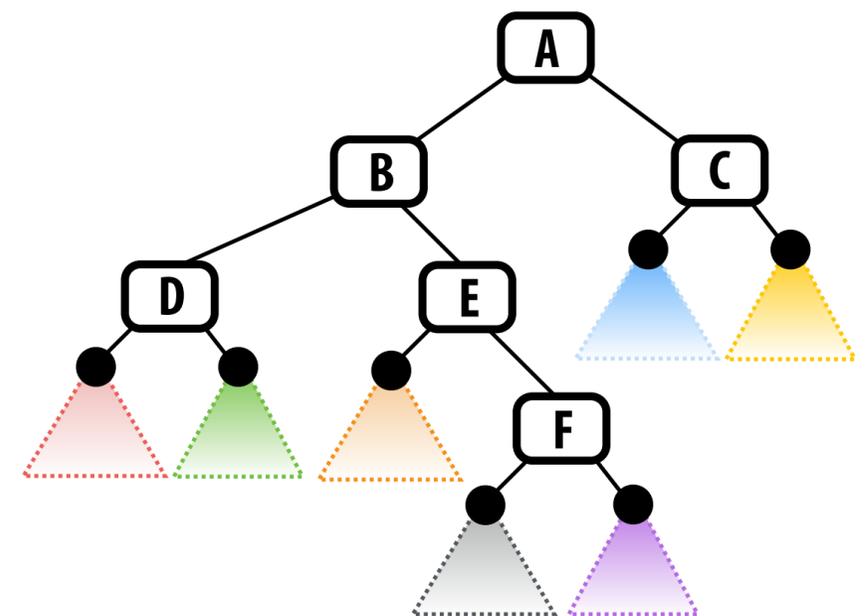
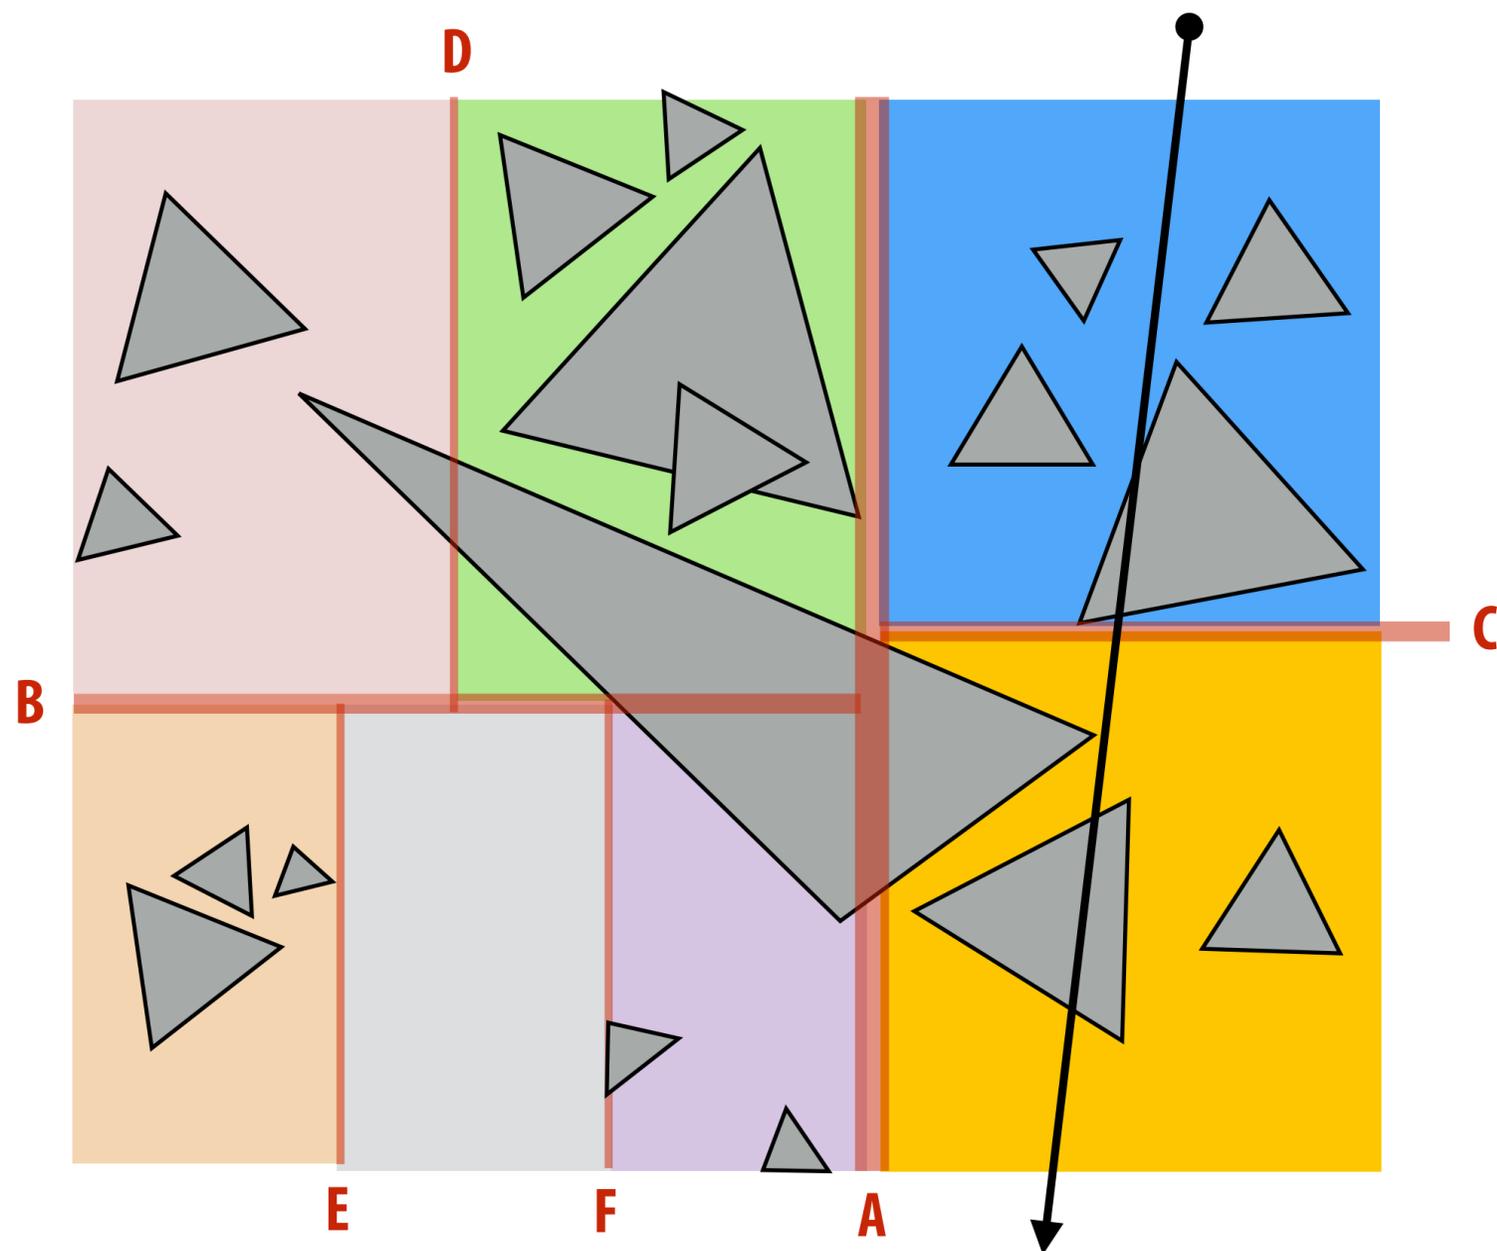
# Primitive-partitioning acceleration structures vs. space-partitioning structures

- **Primitive partitioning (bounding volume hierarchy): partitions primitives into disjoint sets (but sets of primitives may overlap in space)**
- **Space-partitioning (grid, K-D tree) partitions space into disjoint regions (primitives may be contained in multiple regions of space)**



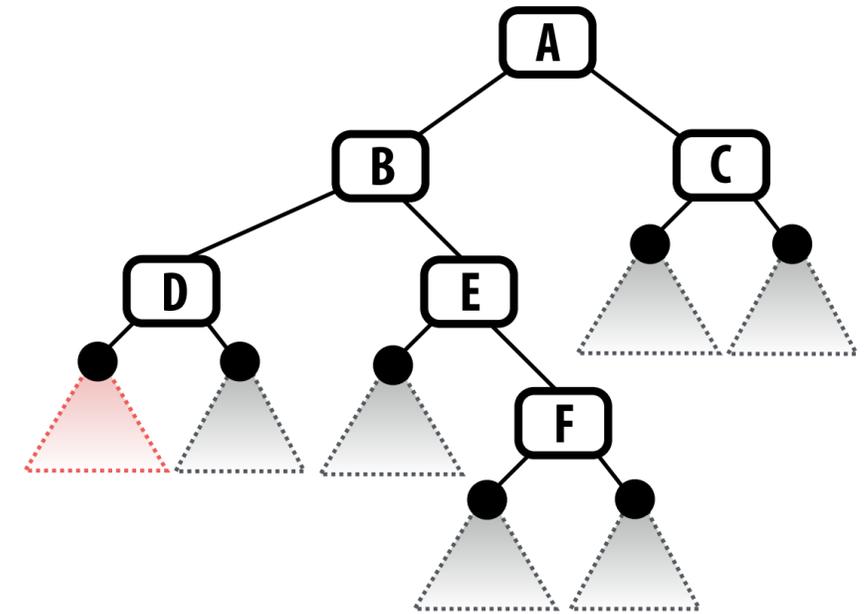
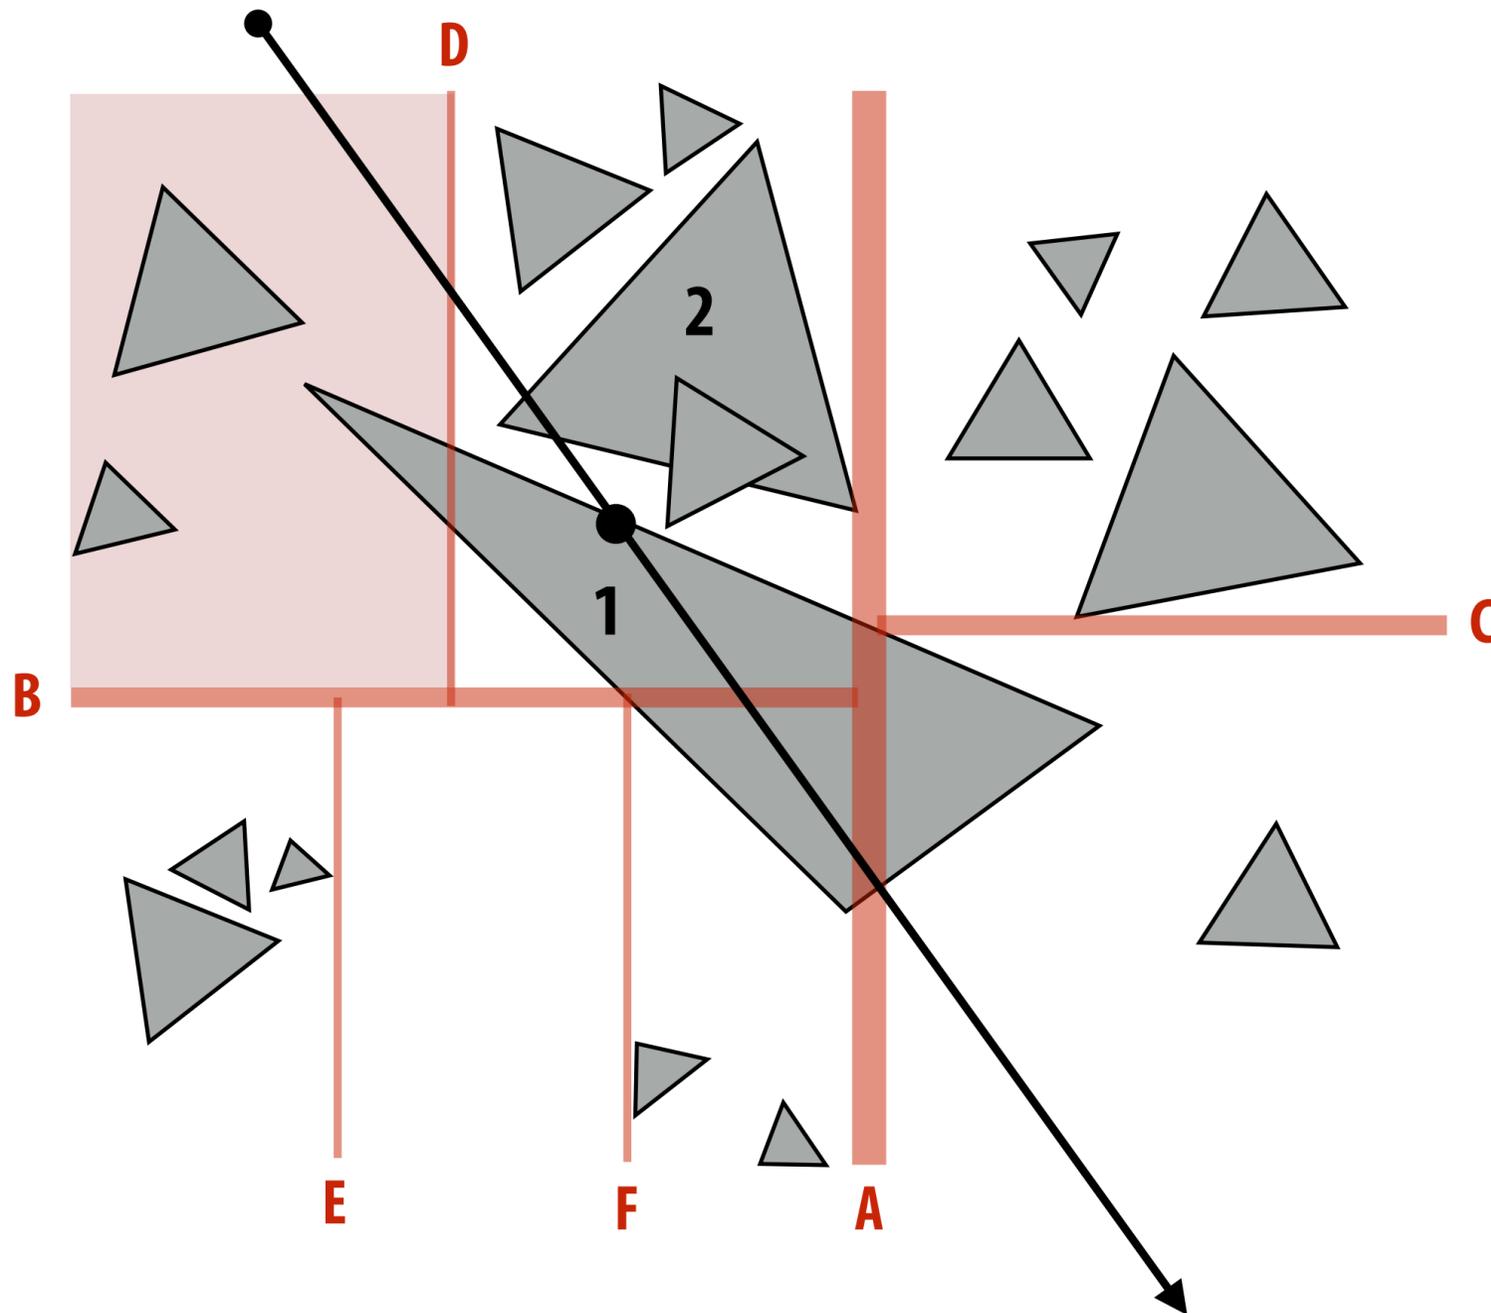
# K-D tree

- **Recursively partition space via axis-aligned partitioning planes**
  - Interior nodes correspond to spatial splits
  - Node traversal can proceed in front-to-back order
  - Unlike BVH, can terminate search after first hit is found.



# Challenge: objects overlap multiple nodes

- Want node traversal to proceed in front-to-back order so traversal can terminate search after first hit found



**Triangle 1 overlaps multiple nodes.**

**Ray hits triangle 1 when in highlighted leaf cell.**

**But intersection with triangle 2 is closer!  
(Haven't traversed to that node yet)**

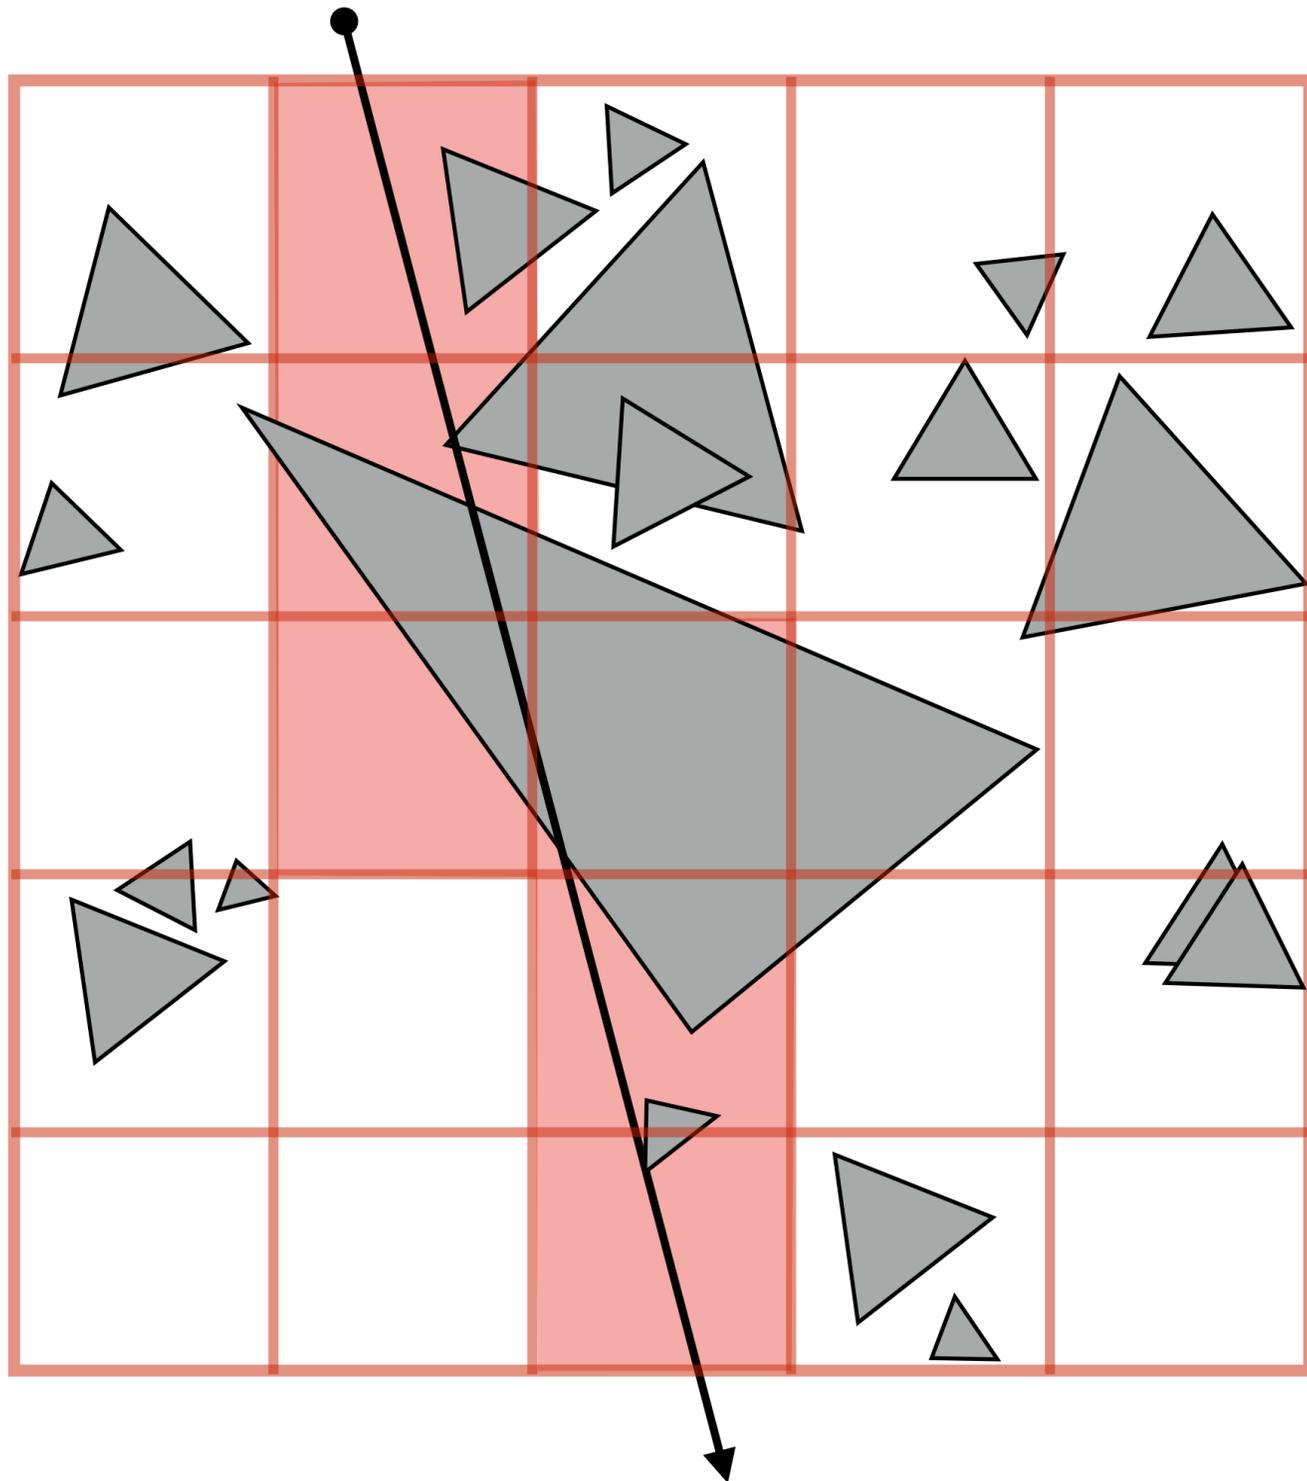
**Solution: require primitive intersection point to be within current leaf node.**

**(primitives may be intersected multiple times by same ray \*)**

\* Caching hit info or "mailboxing" can be used to avoid repeated intersections

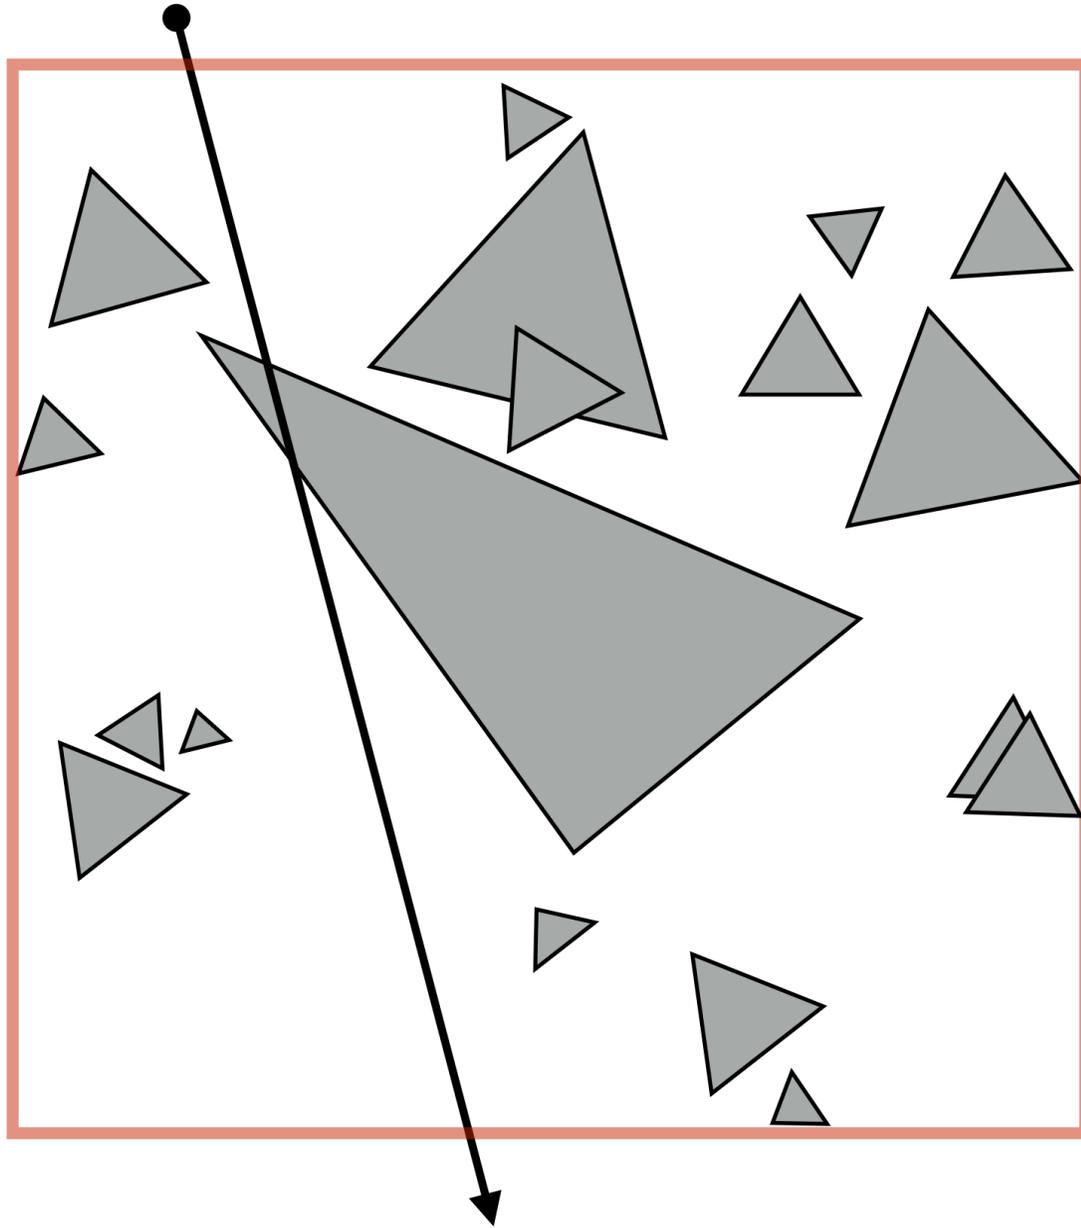
# **Uniform grid (a very simple hierarchy)**

# Uniform grid

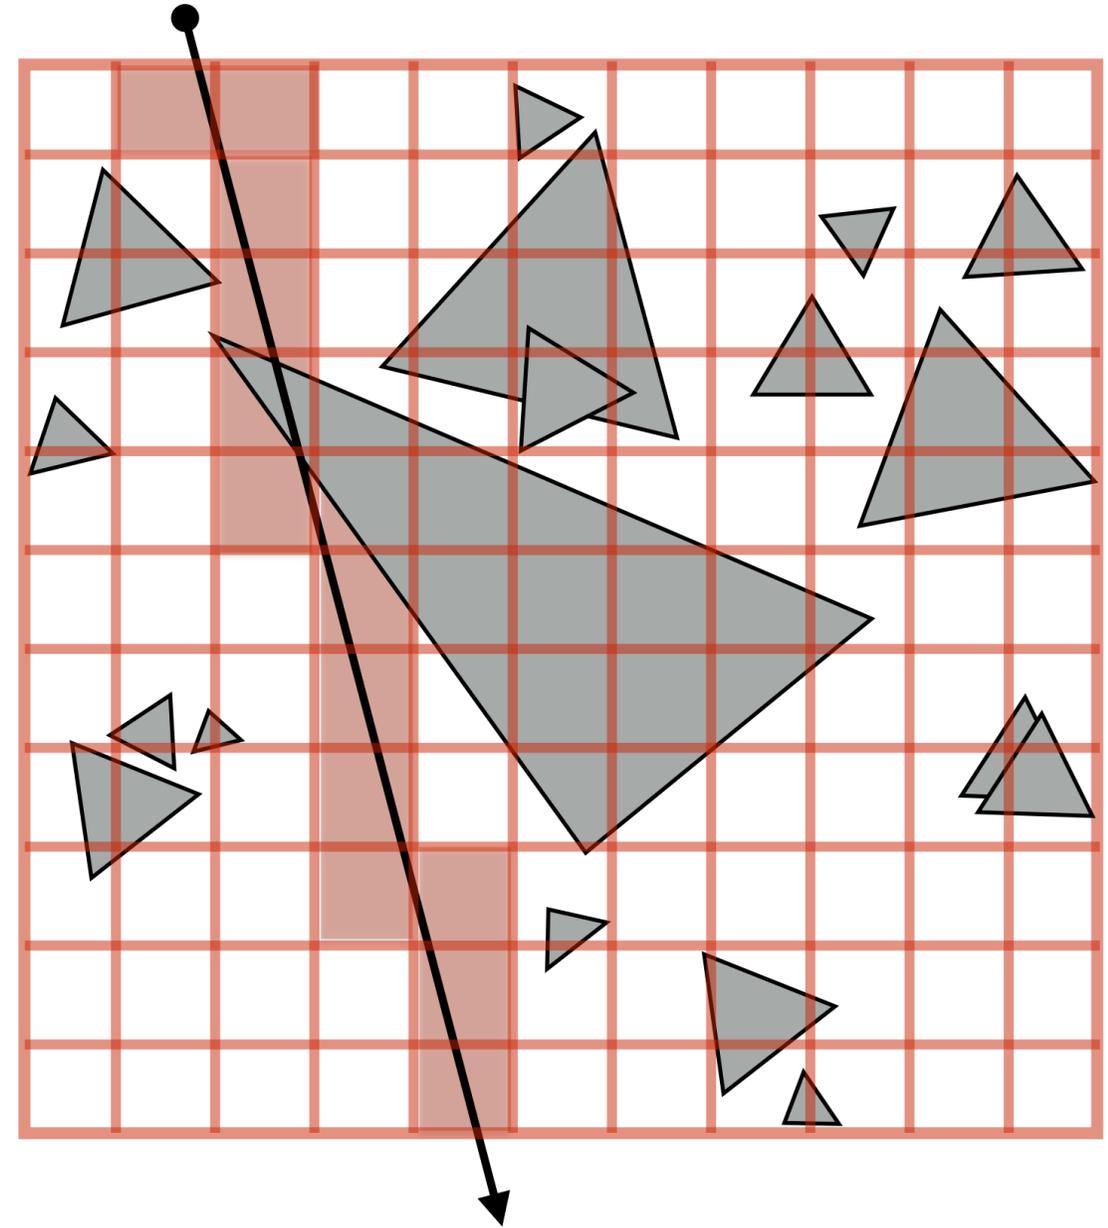


- Partition space into equal sized volumes (volume-elements or “voxels”)
- Each grid cell contains primitives that overlap the voxel. (very cheap to construct acceleration structure)
- Walk ray through volume in order
  - Very efficient implementation possible (think: *3D line rasterization*)
  - Only consider intersection with primitives in voxels the ray intersects

# What should the grid resolution be?



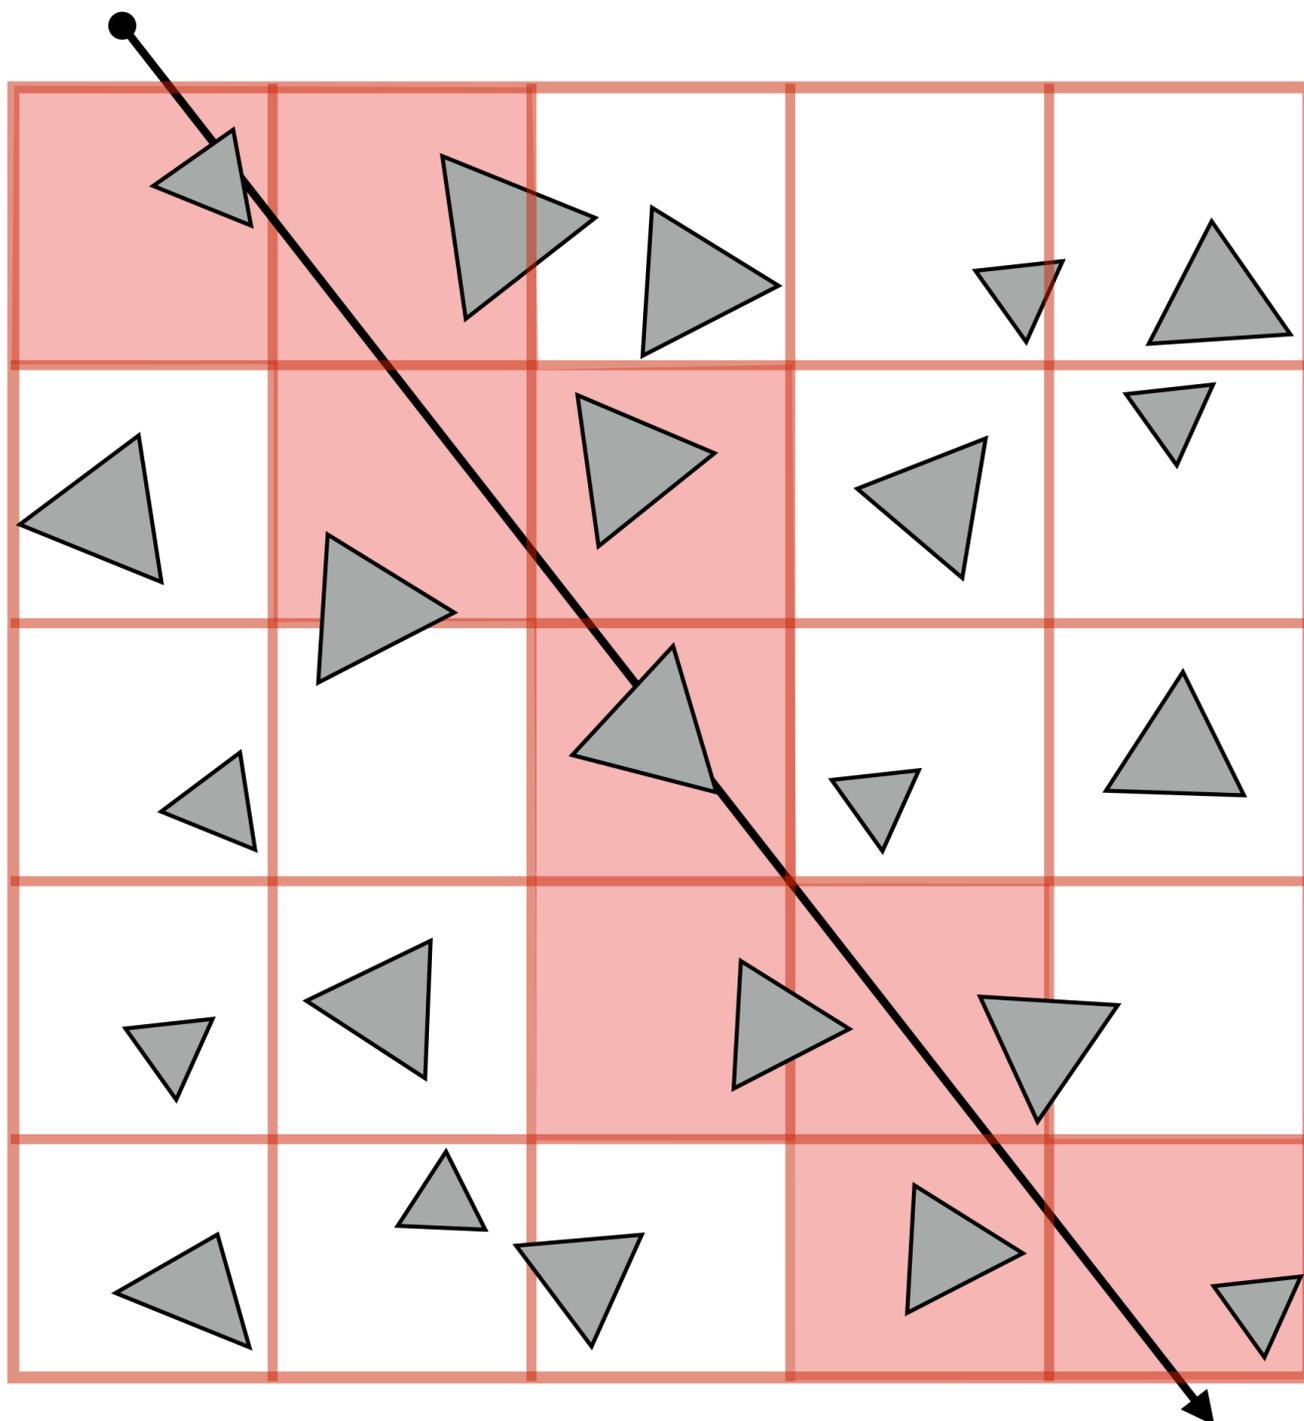
**Too few grids cell: degenerates to brute-force approach**



**Too many grid cells: incur significant cost traversing through cells with empty space**

# Heuristic

- **Choose number of voxels  $\sim$  total number of primitives**  
(constant prims per voxel — assuming uniform distribution of primitives)



Intersection cost:  $O(\sqrt[3]{N})$   
(assuming 3D grid)

**(Q: Which grows faster,  
cube root of N or log(N)?**

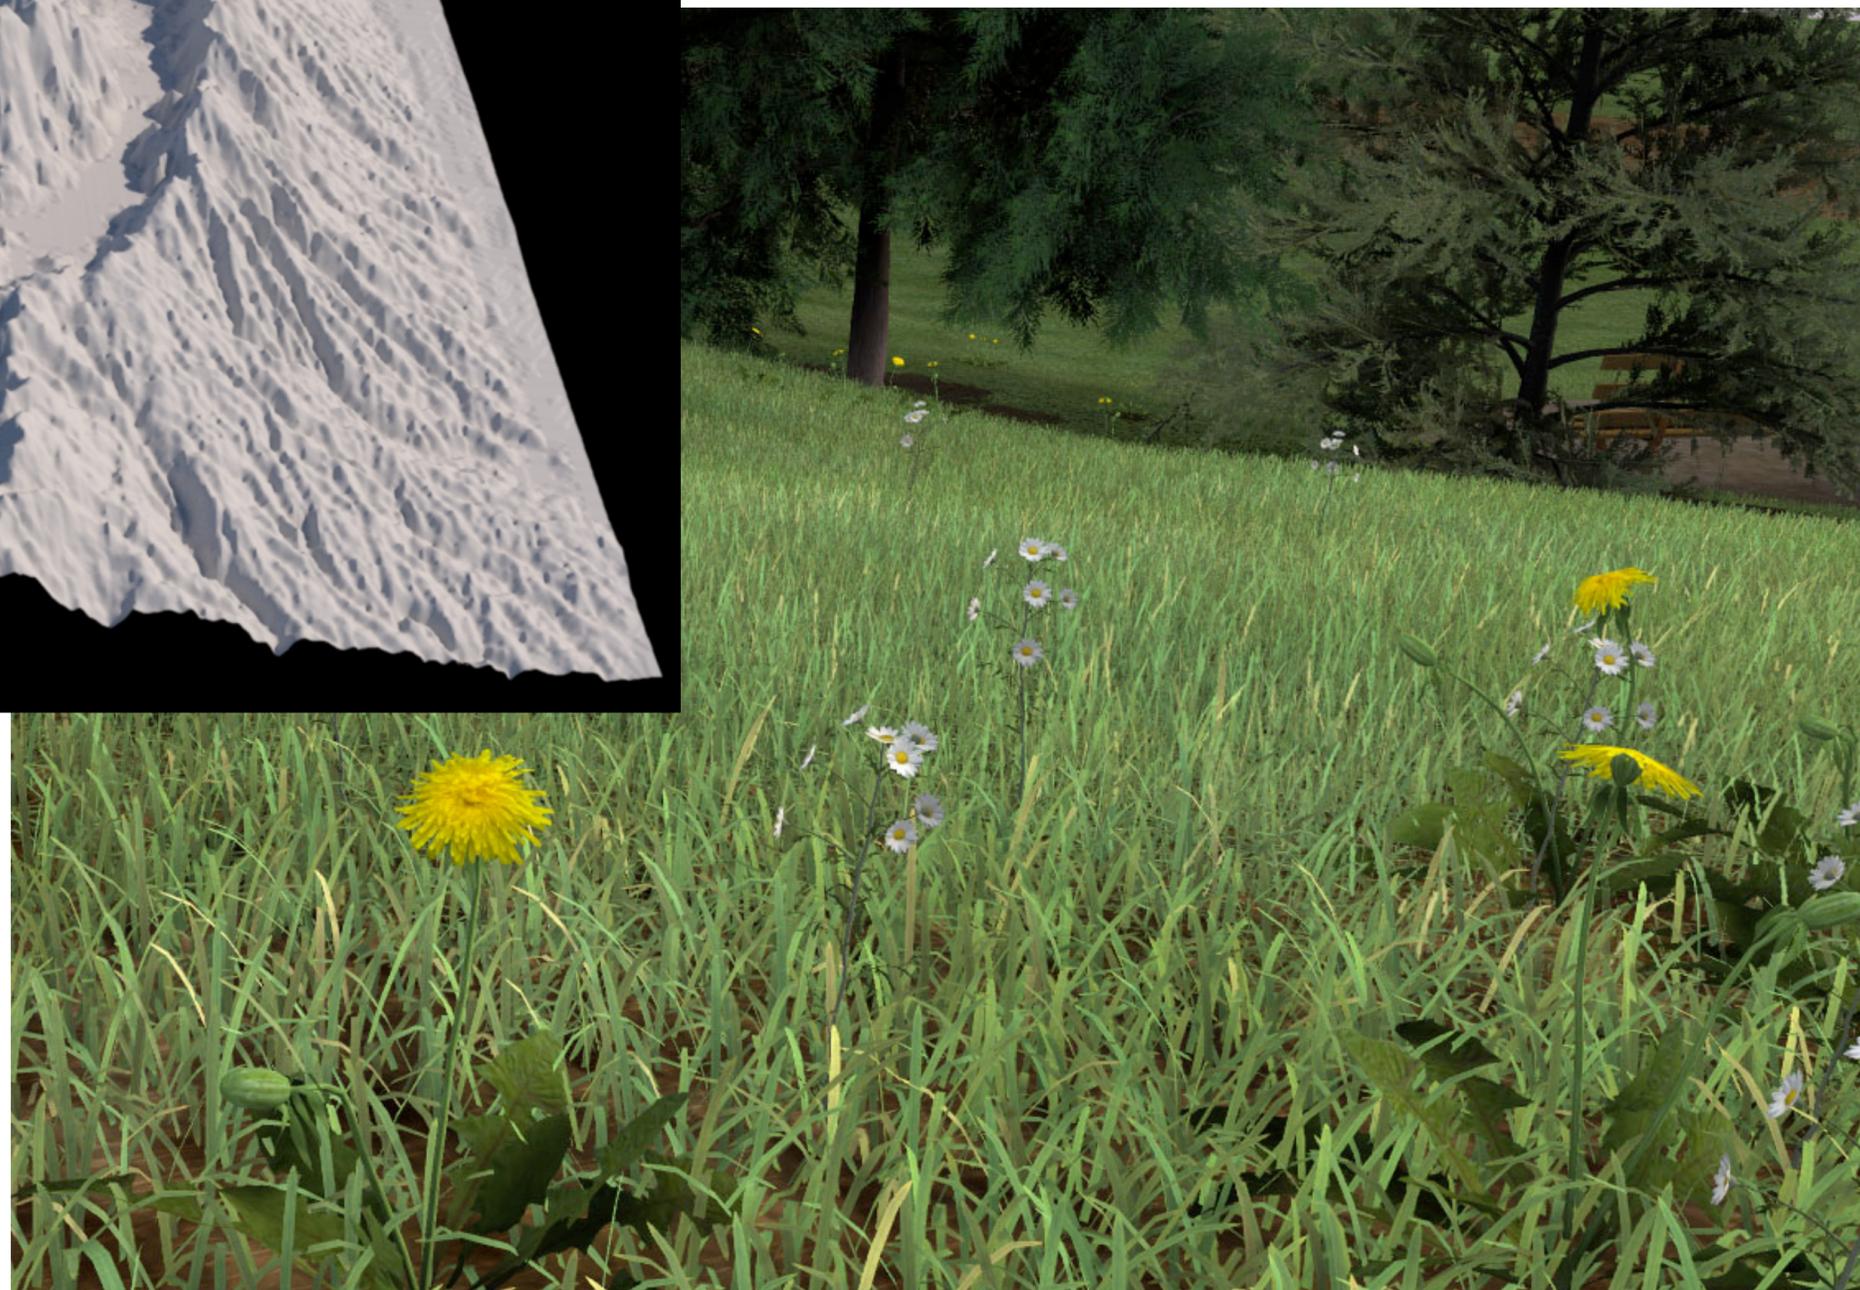
# When uniform grids work well: uniform distribution of primitives in scene



**Terrain / height fields:**

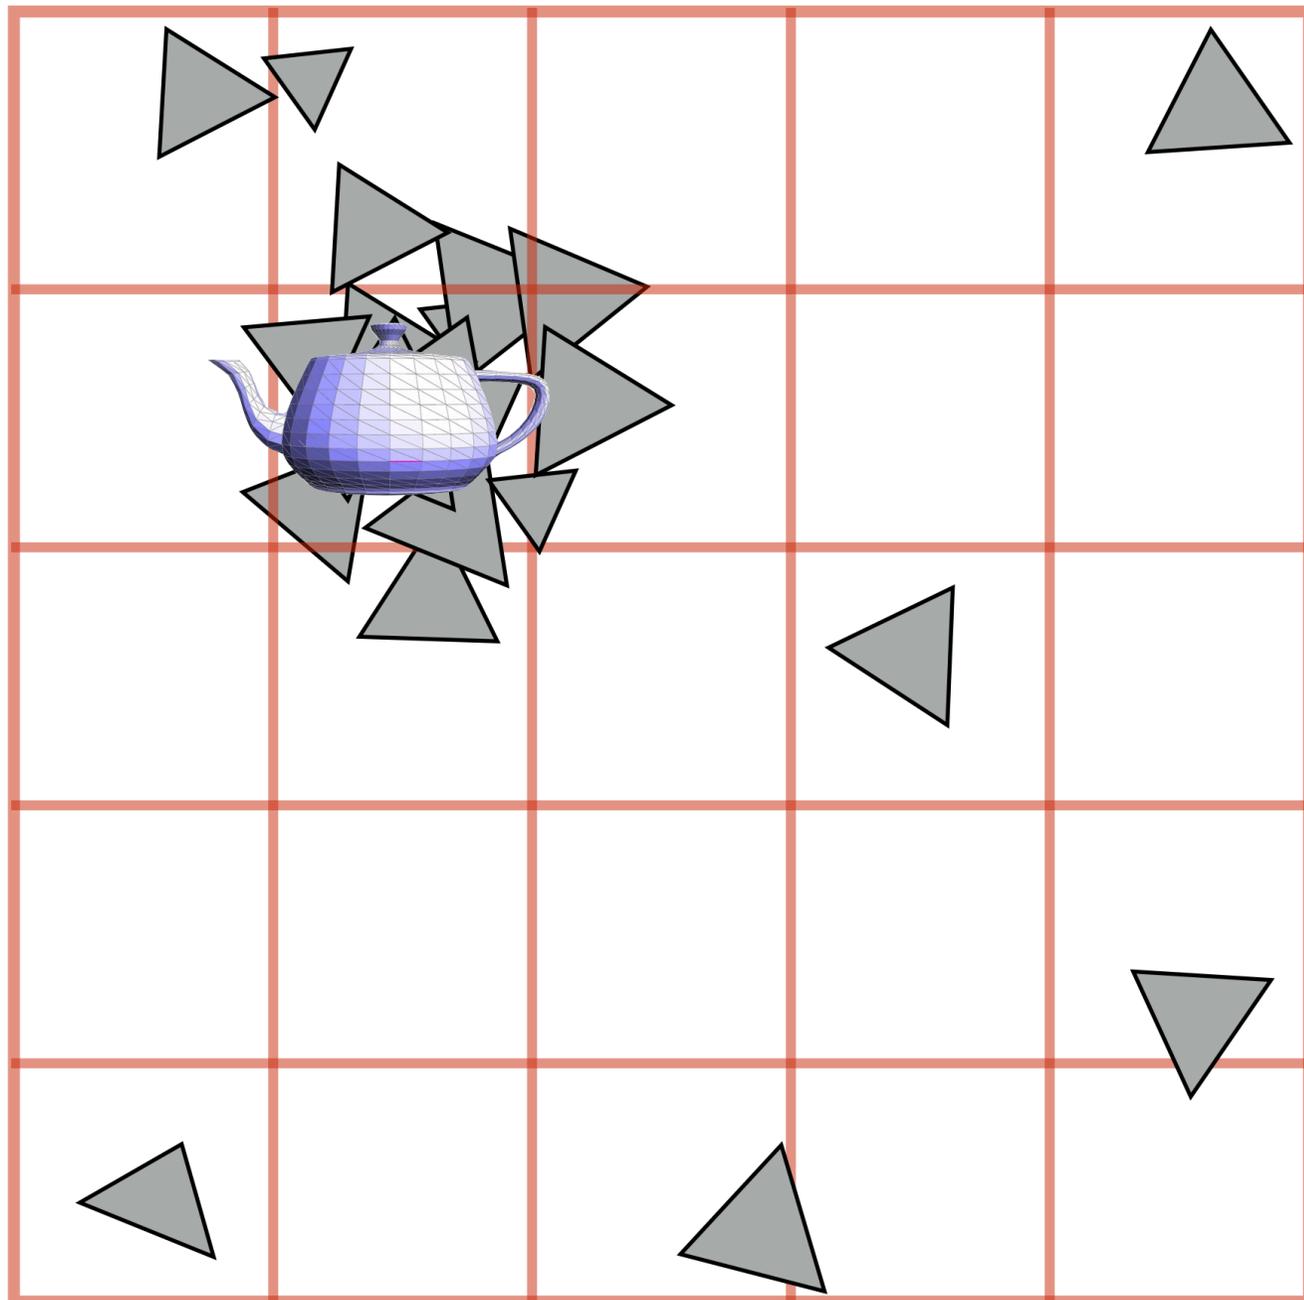
[Image credit: Misuba Renderer]

**Grass:**



[Image credit: [www.kevinboulanger.net/grass.html](http://www.kevinboulanger.net/grass.html)]

# Uniform grids cannot adapt to non-uniform distribution of geometry in scene



**“Teapot in a stadium problem”**

**Scene has large spatial extent.**

**Contains a high-resolution object that has small spatial extent (ends up in one grid cell)**

# When uniform grids do not work well: non-uniform distribution of geometric detail



Jun Yan, Tracy Renderer

# When uniform grids do not work well: non-uniform distribution of geometric detail

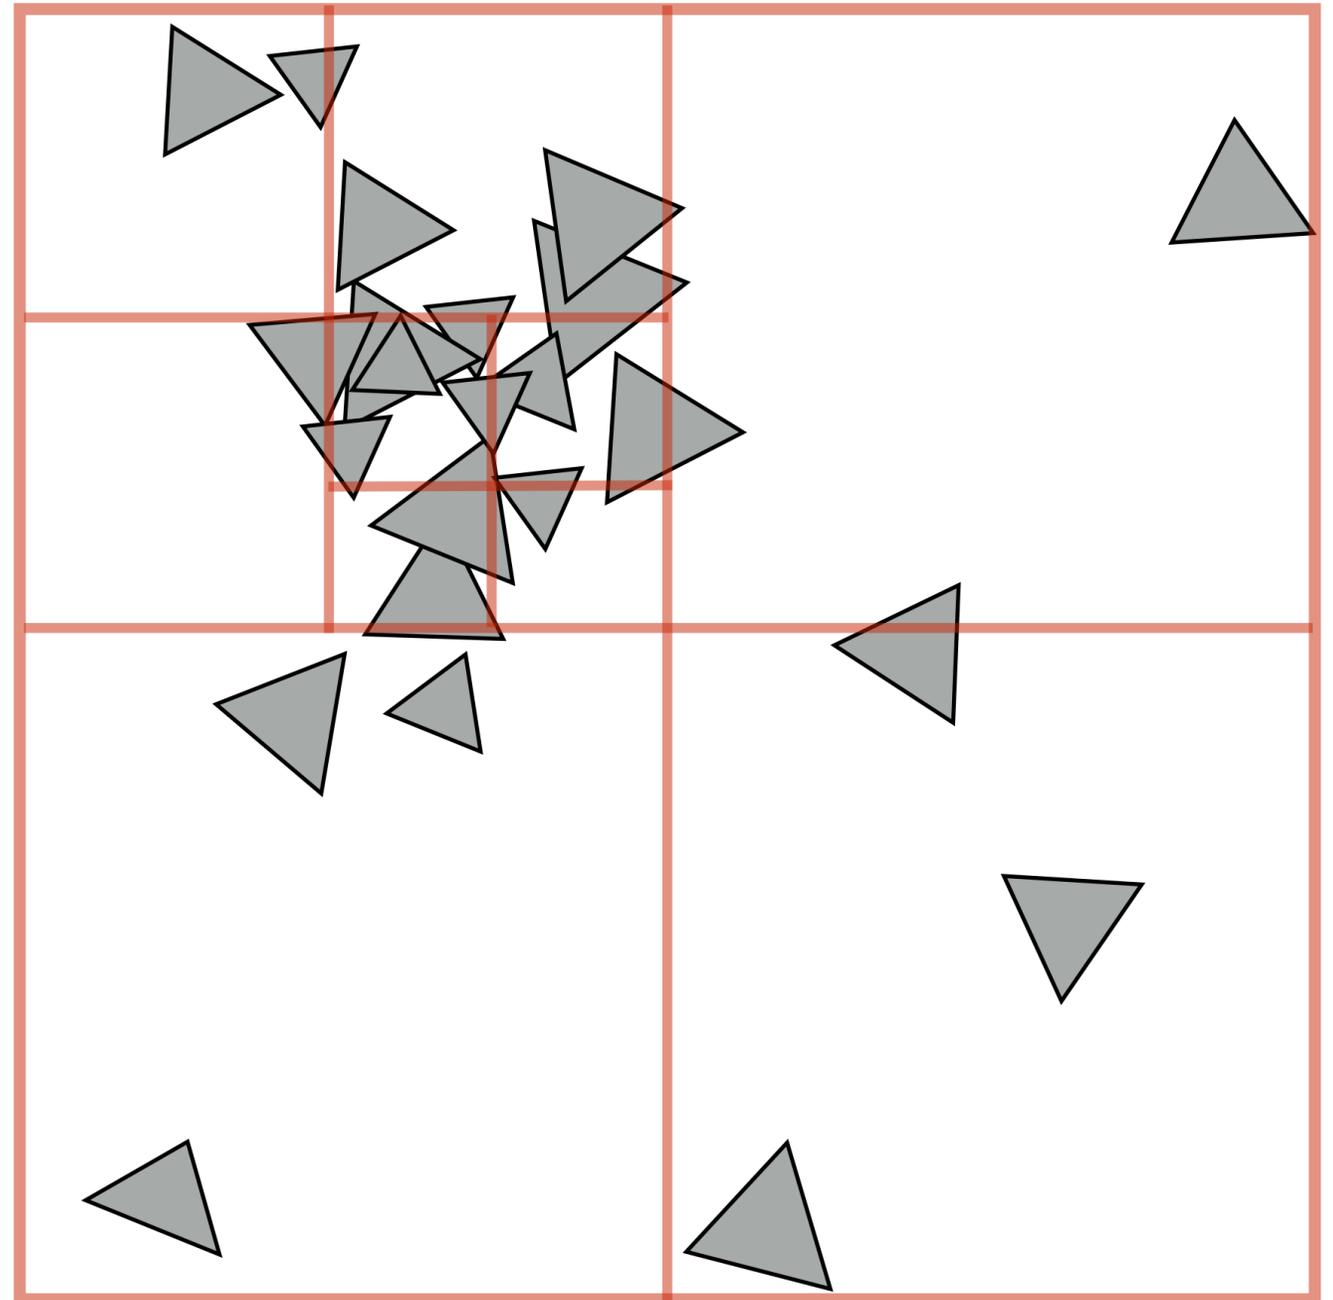


# Quad-tree / octree

**Like uniform grid: easy to build (don't have to choose partition planes)**

**Has greater ability to adapt to location of scene geometry than uniform grid.**

**But lower intersection performance than K-D tree (only limited ability to adapt)**



**Quad-tree: nodes have 4 children (partitions 2D space)**

**Octree: nodes have 8 children (partitions 3D space)**

# Summary of spatial acceleration structures:

## *Choose the right structure for the job!*

- **Primitive vs. spatial partitioning:**
  - **Primitive partitioning: partition sets of objects**
    - Bounded number of BVH nodes, *simpler to update if primitives in scene change position*
  - **Spatial partitioning: partition space**
    - Traverse space in order (first intersection is closest intersection), may intersect primitive multiple times
- **Adaptive structures (BVH, K-D tree)**
  - **More costly to construct (must be able to amortize cost over many geometric queries)**
  - **Better intersection performance under non-uniform distribution of primitives**
- **Non-adaptive accelerations structures (uniform grids)**
  - **Simple, cheap to construct**
  - **Good intersection performance if scene primitives are uniformly distributed**
- **Many, many combinations thereof...**

**Rendering via ray casting:  
one common use of ray-scene intersection tests**

**Rasterization and ray casting are two algorithms for solving the same problem: determining “visibility from a camera”**

# Recall triangle visibility:

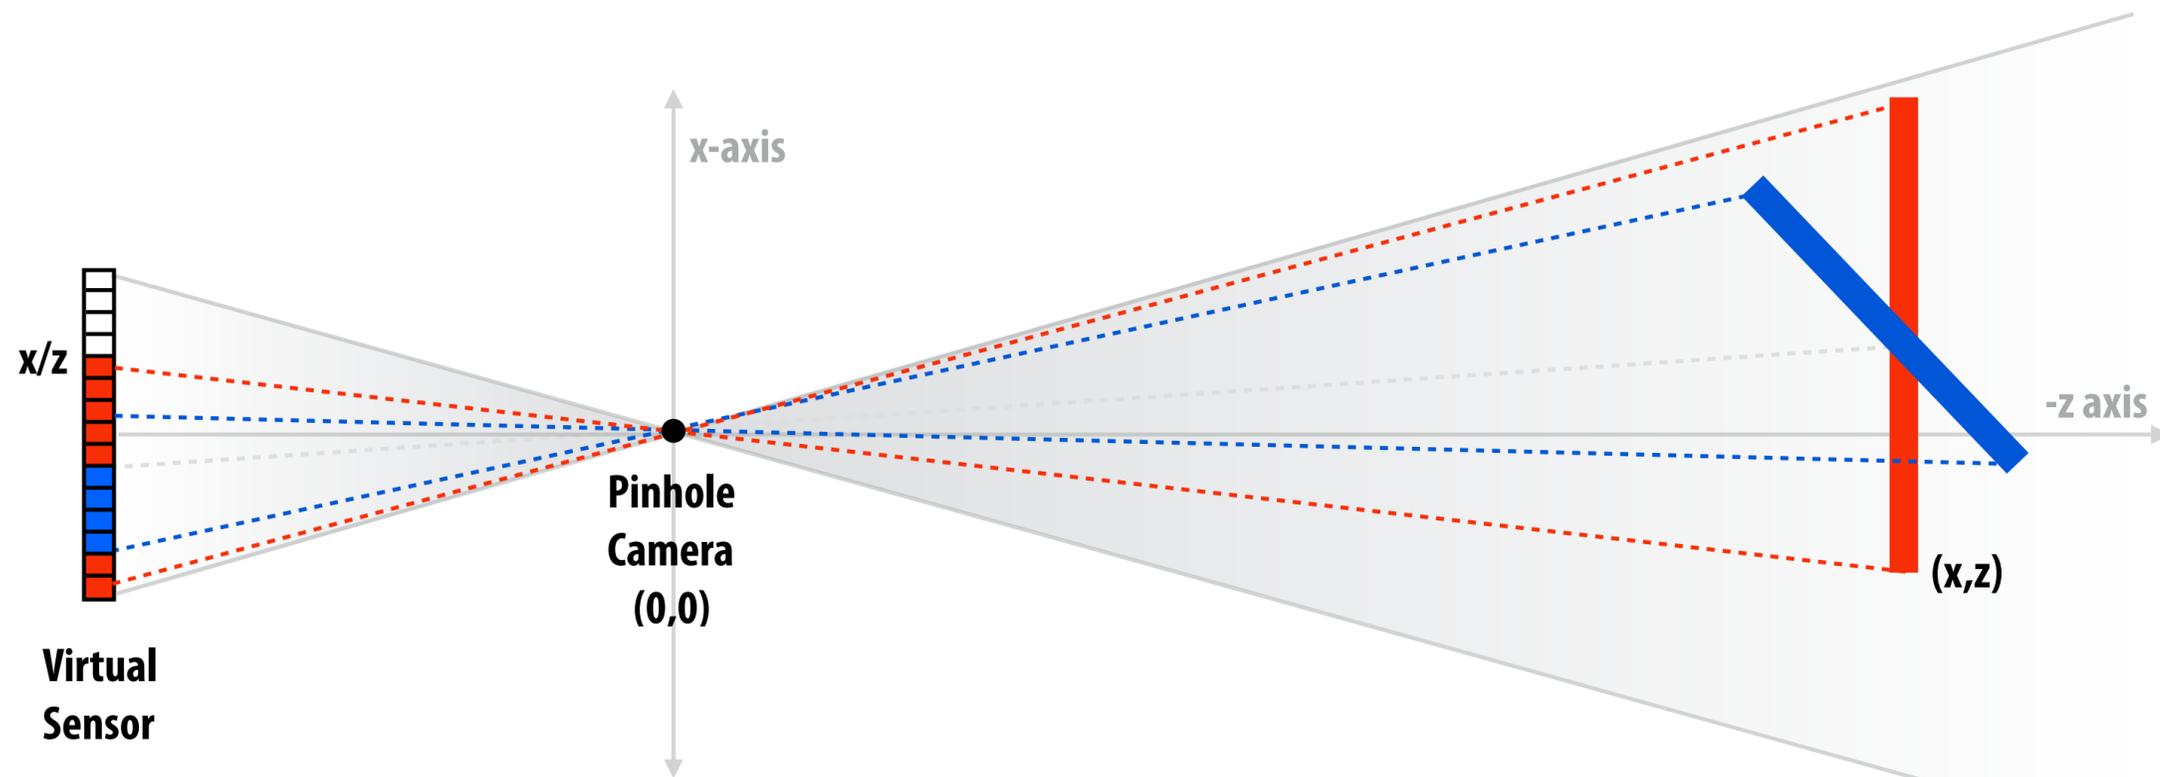
**Question 1: what samples does the triangle overlap?  
("coverage")**

**Sample**

**Question 2: what triangle is closest to the  
camera in each sample? ("occlusion")**

# The visibility problem

- **What scene geometry is visible at each screen sample?**
  - What scene geometry *projects* onto screen sample points? (coverage)
  - Which geometry is visible from the camera at each sample? (occlusion)



# Basic rasterization algorithm

Sample = 2D point

Coverage: 2D triangle/sample tests (does projected triangle cover 2D sample point)

Occlusion: depth buffer

```
initialize z_closest[] to INFINITY // store closest-surface-so-far for all samples
initialize color[] // store scene color for all samples
for each triangle t in scene: // loop 1: triangles
    t_proj = project_triangle(t)
    for each 2D sample s in frame buffer: // loop 2: visibility samples
        if (t_proj covers s)
            compute color of triangle at sample
            if (depth of t at s is closer than z_closest[s])
                update z_closest[s] and color[s]
```

***“Given a triangle, find the samples it covers”***

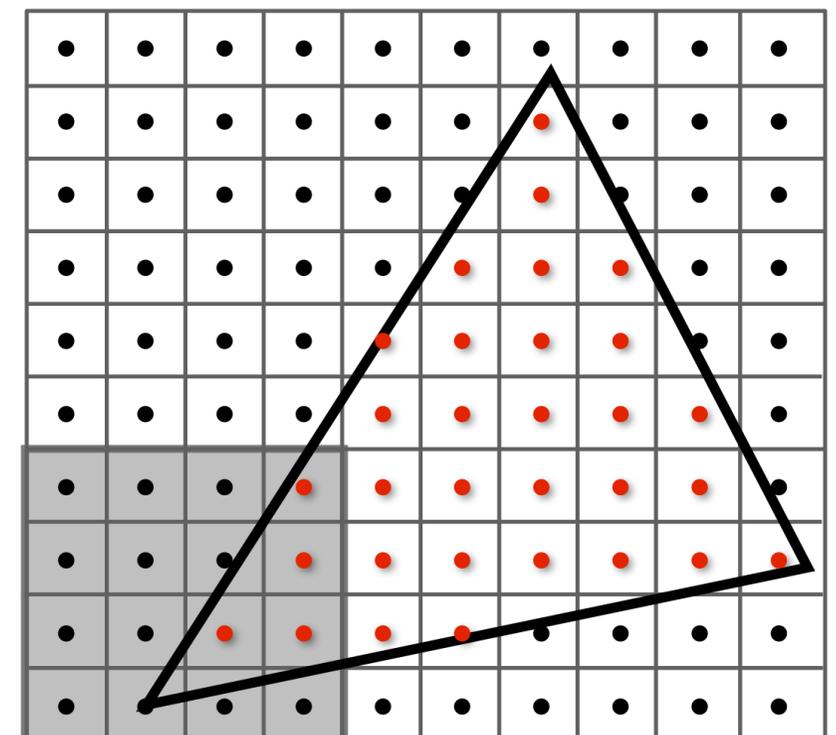
(finding the samples is relatively easy since they are distributed uniformly on screen)

More modern hierarchical rasterization:

For each TILE of image

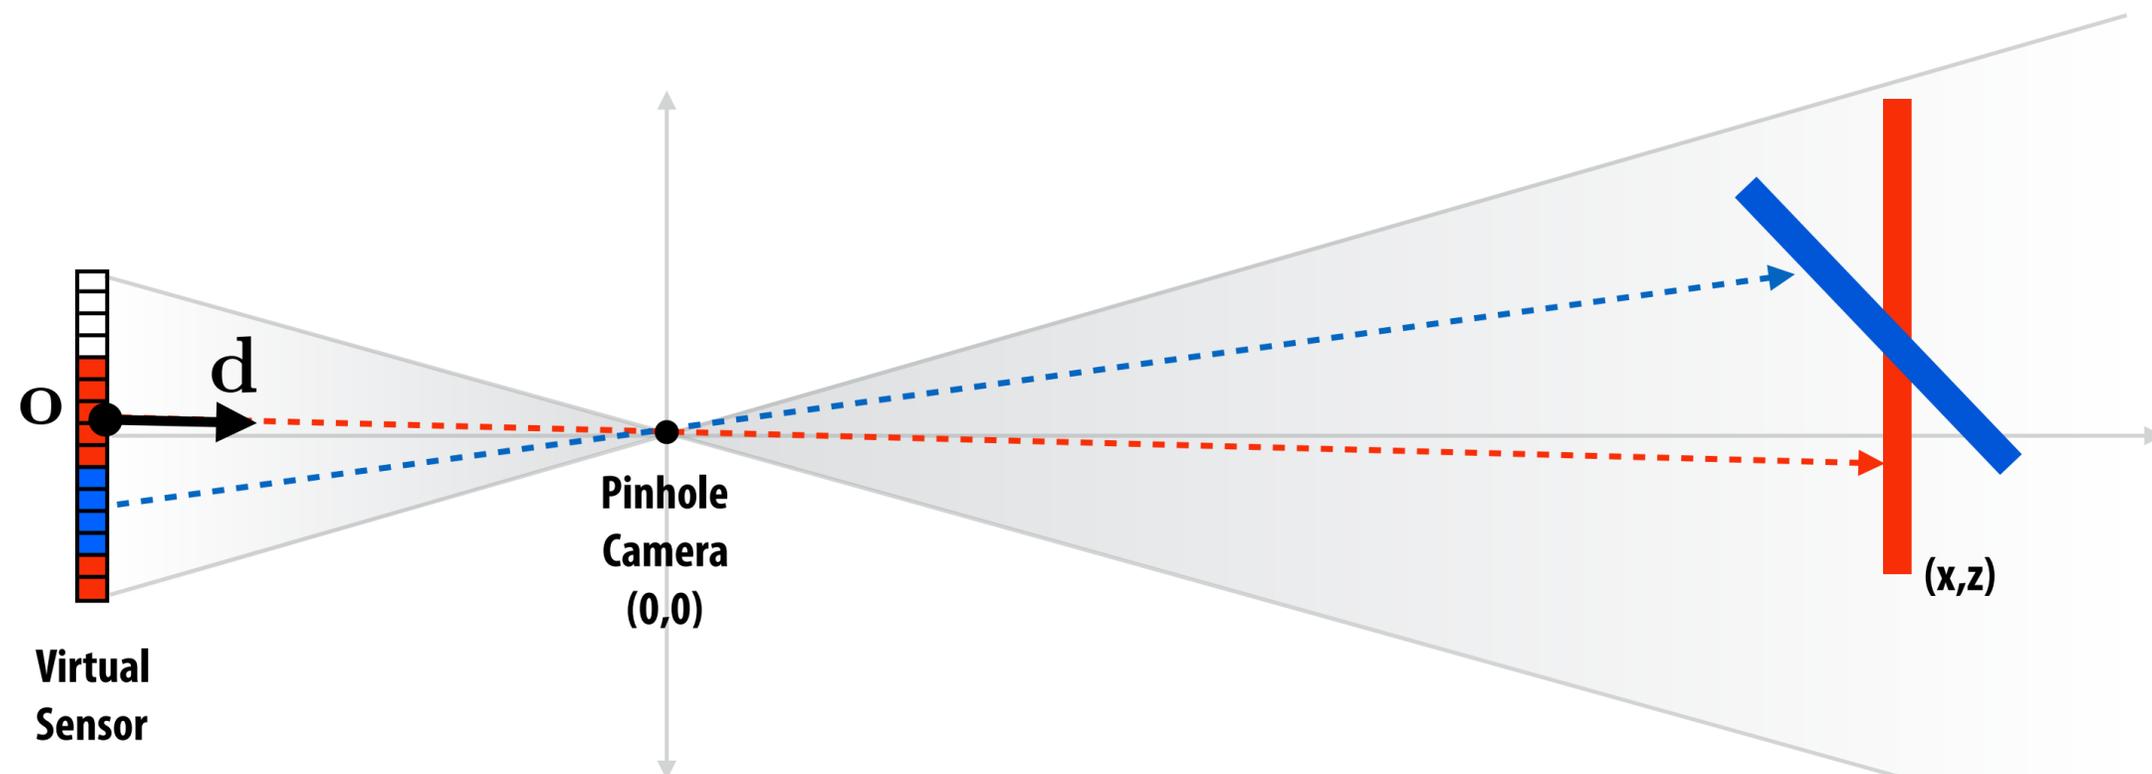
If triangle overlaps tile, check all samples in tile

*(What does this strategy remind you of? :-))*



# The visibility problem (described differently)

- In terms of casting rays from the camera:
  - Is a scene primitive hit by a ray originating from a point on the virtual sensor and traveling through the aperture of the pinhole camera? (coverage)
  - What primitive is the first hit along that ray? (occlusion)



# Basic ray casting algorithm

Sample = a ray in 3D

Coverage: 3D ray-triangle intersection tests (does ray “hit” triangle)

Occlusion: closest intersection along ray

```
initialize color[] // store scene color for all samples
for each sample s in frame buffer: // loop 1: visibility samples (rays)
    r = ray from s on sensor through pinhole aperture
    r.min_t = INFINITY // only store closest-so-far for current ray
    r.tri = NULL;
    for each triangle tri in scene: // loop 2: triangles
        if (intersects(r, tri)) { // 3D ray-triangle intersection test
            if (intersection distance along ray is closer than r.min_t)
                update r.min_t and r.tri = tri;
        }
    color[s] = compute surface color of triangle r.tri at hit point
```

Compared to rasterization approach: just a reordering of the loops!

*“Given a ray, find the closest triangle it hits.”*

As we saw today, the brute force “for each triangle” loop above is typically accelerated using an acceleration structure. (A rasterizer’s “for each sample” inner loop is not just a loop over all screen samples either!)

# Basic rasterization vs. ray casting

## ■ Rasterization:

- Proceeds in triangle order
- Store entire depth buffer (random access to regular structure of fixed size)
- Don't have to store entire scene geometry in memory, naturally supports unbounded size scenes

## ■ Ray casting:

- Proceeds in screen sample order
  - Don't have to store closest depth so far for the entire screen (just current ray)
  - Natural order for rendering transparent surfaces (process surfaces in the order the are encountered along the ray: front-to-back or back-to-front)
- Must store entire scene geometry
- Performance more strongly depends on distribution of primitives in scene

## ■ Modern high-performance implementations of rasterization and ray-casting embody very similar techniques

- Hierarchies of rays/samples
- Hierarchies of geometry
- ...

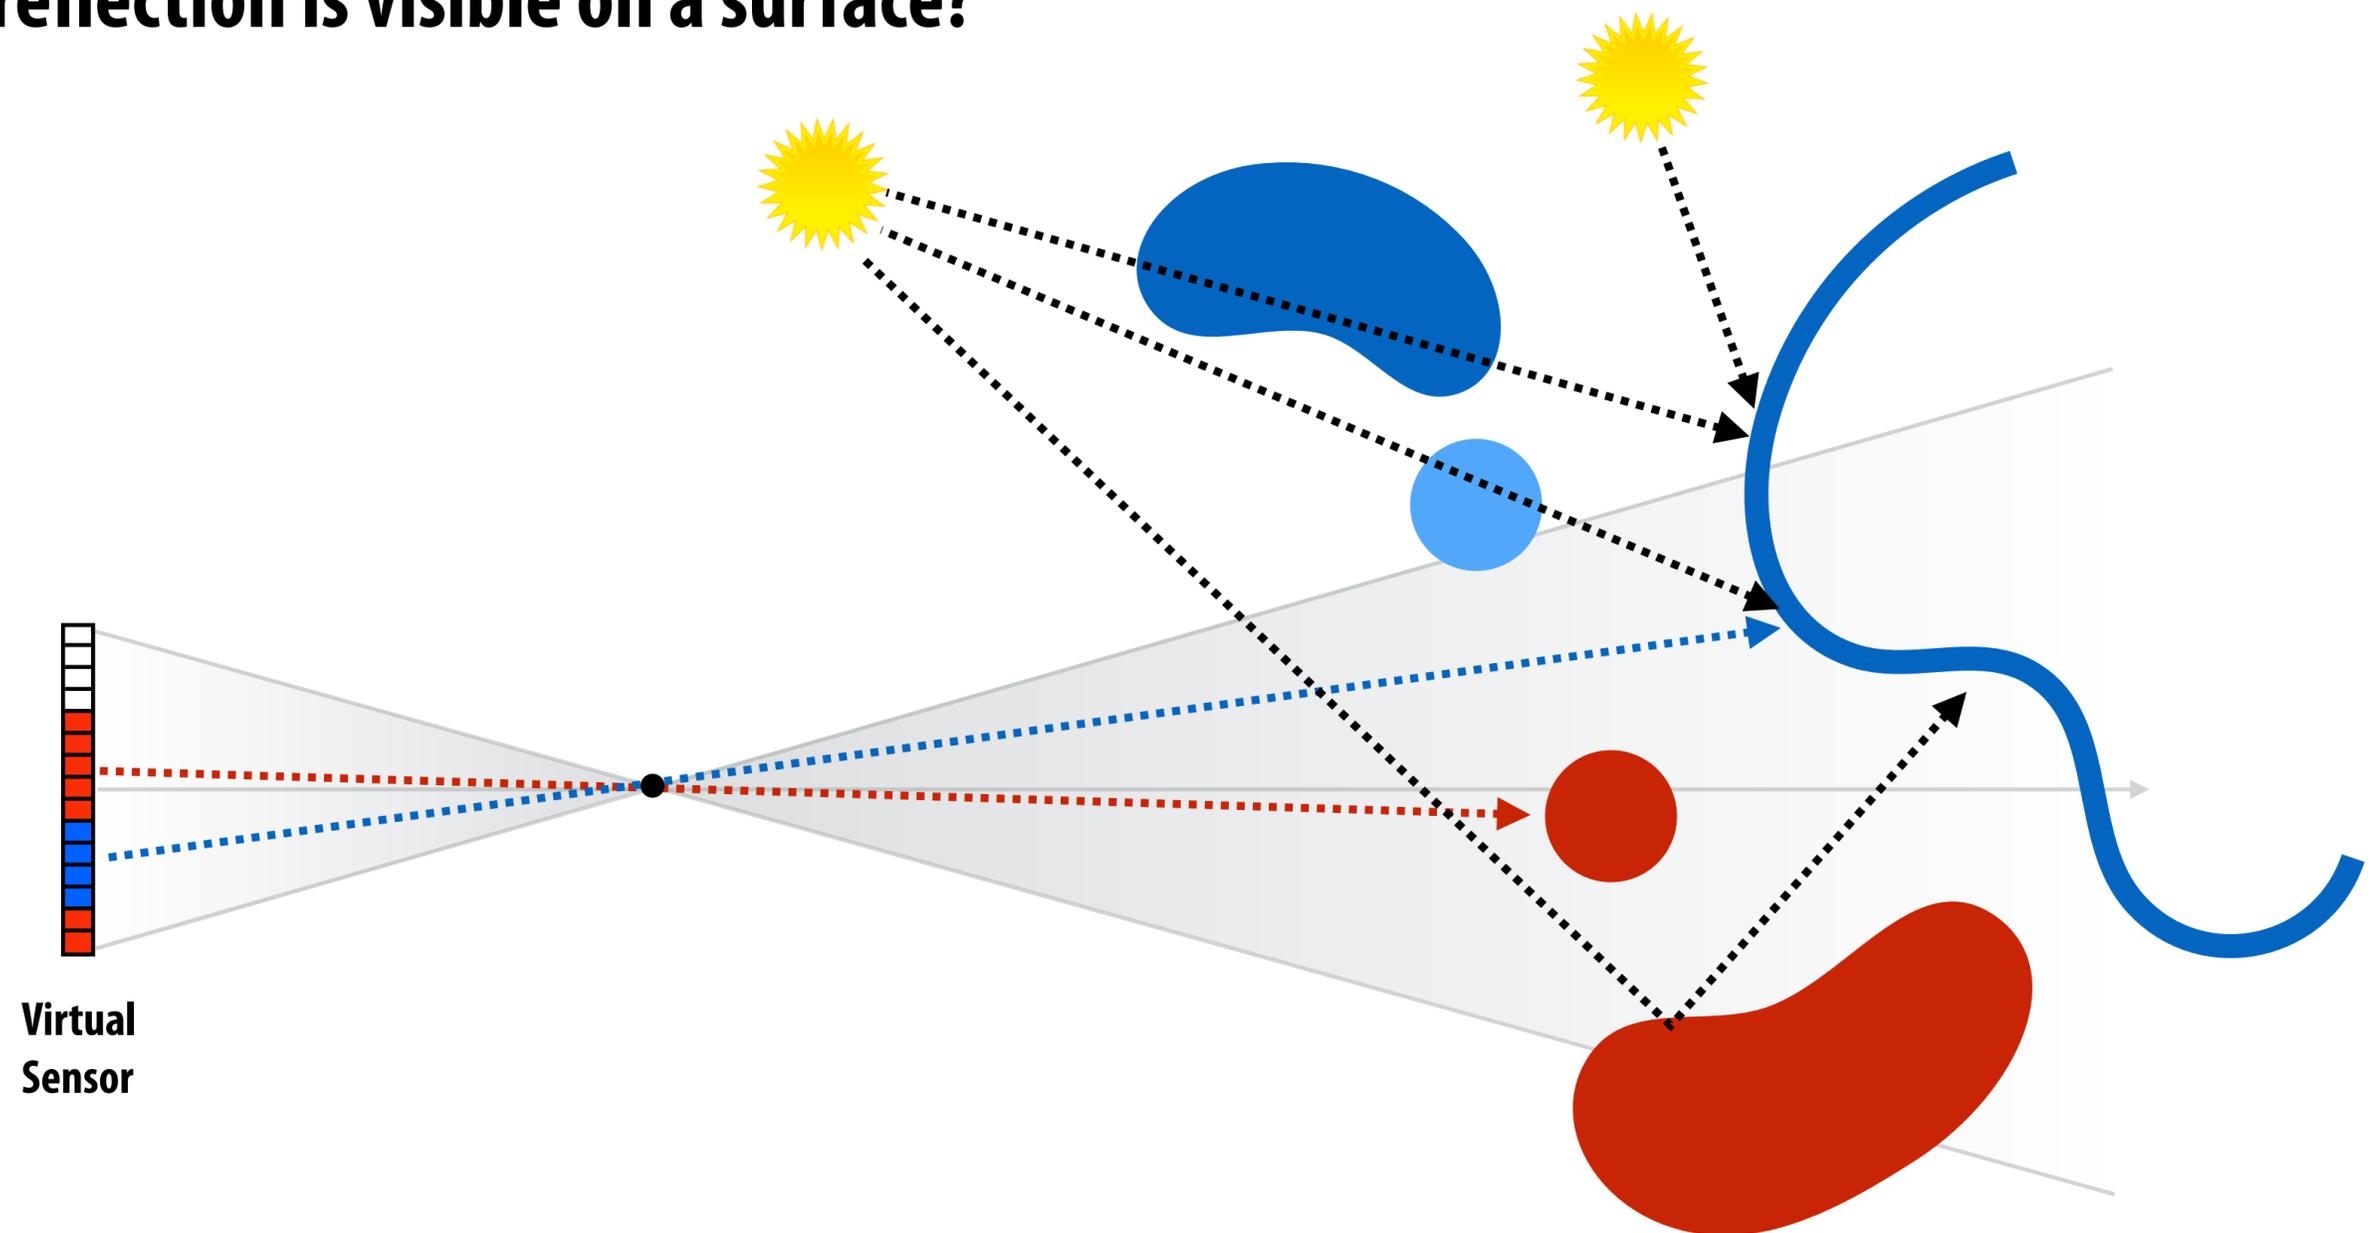


# Generality of ray-scene queries

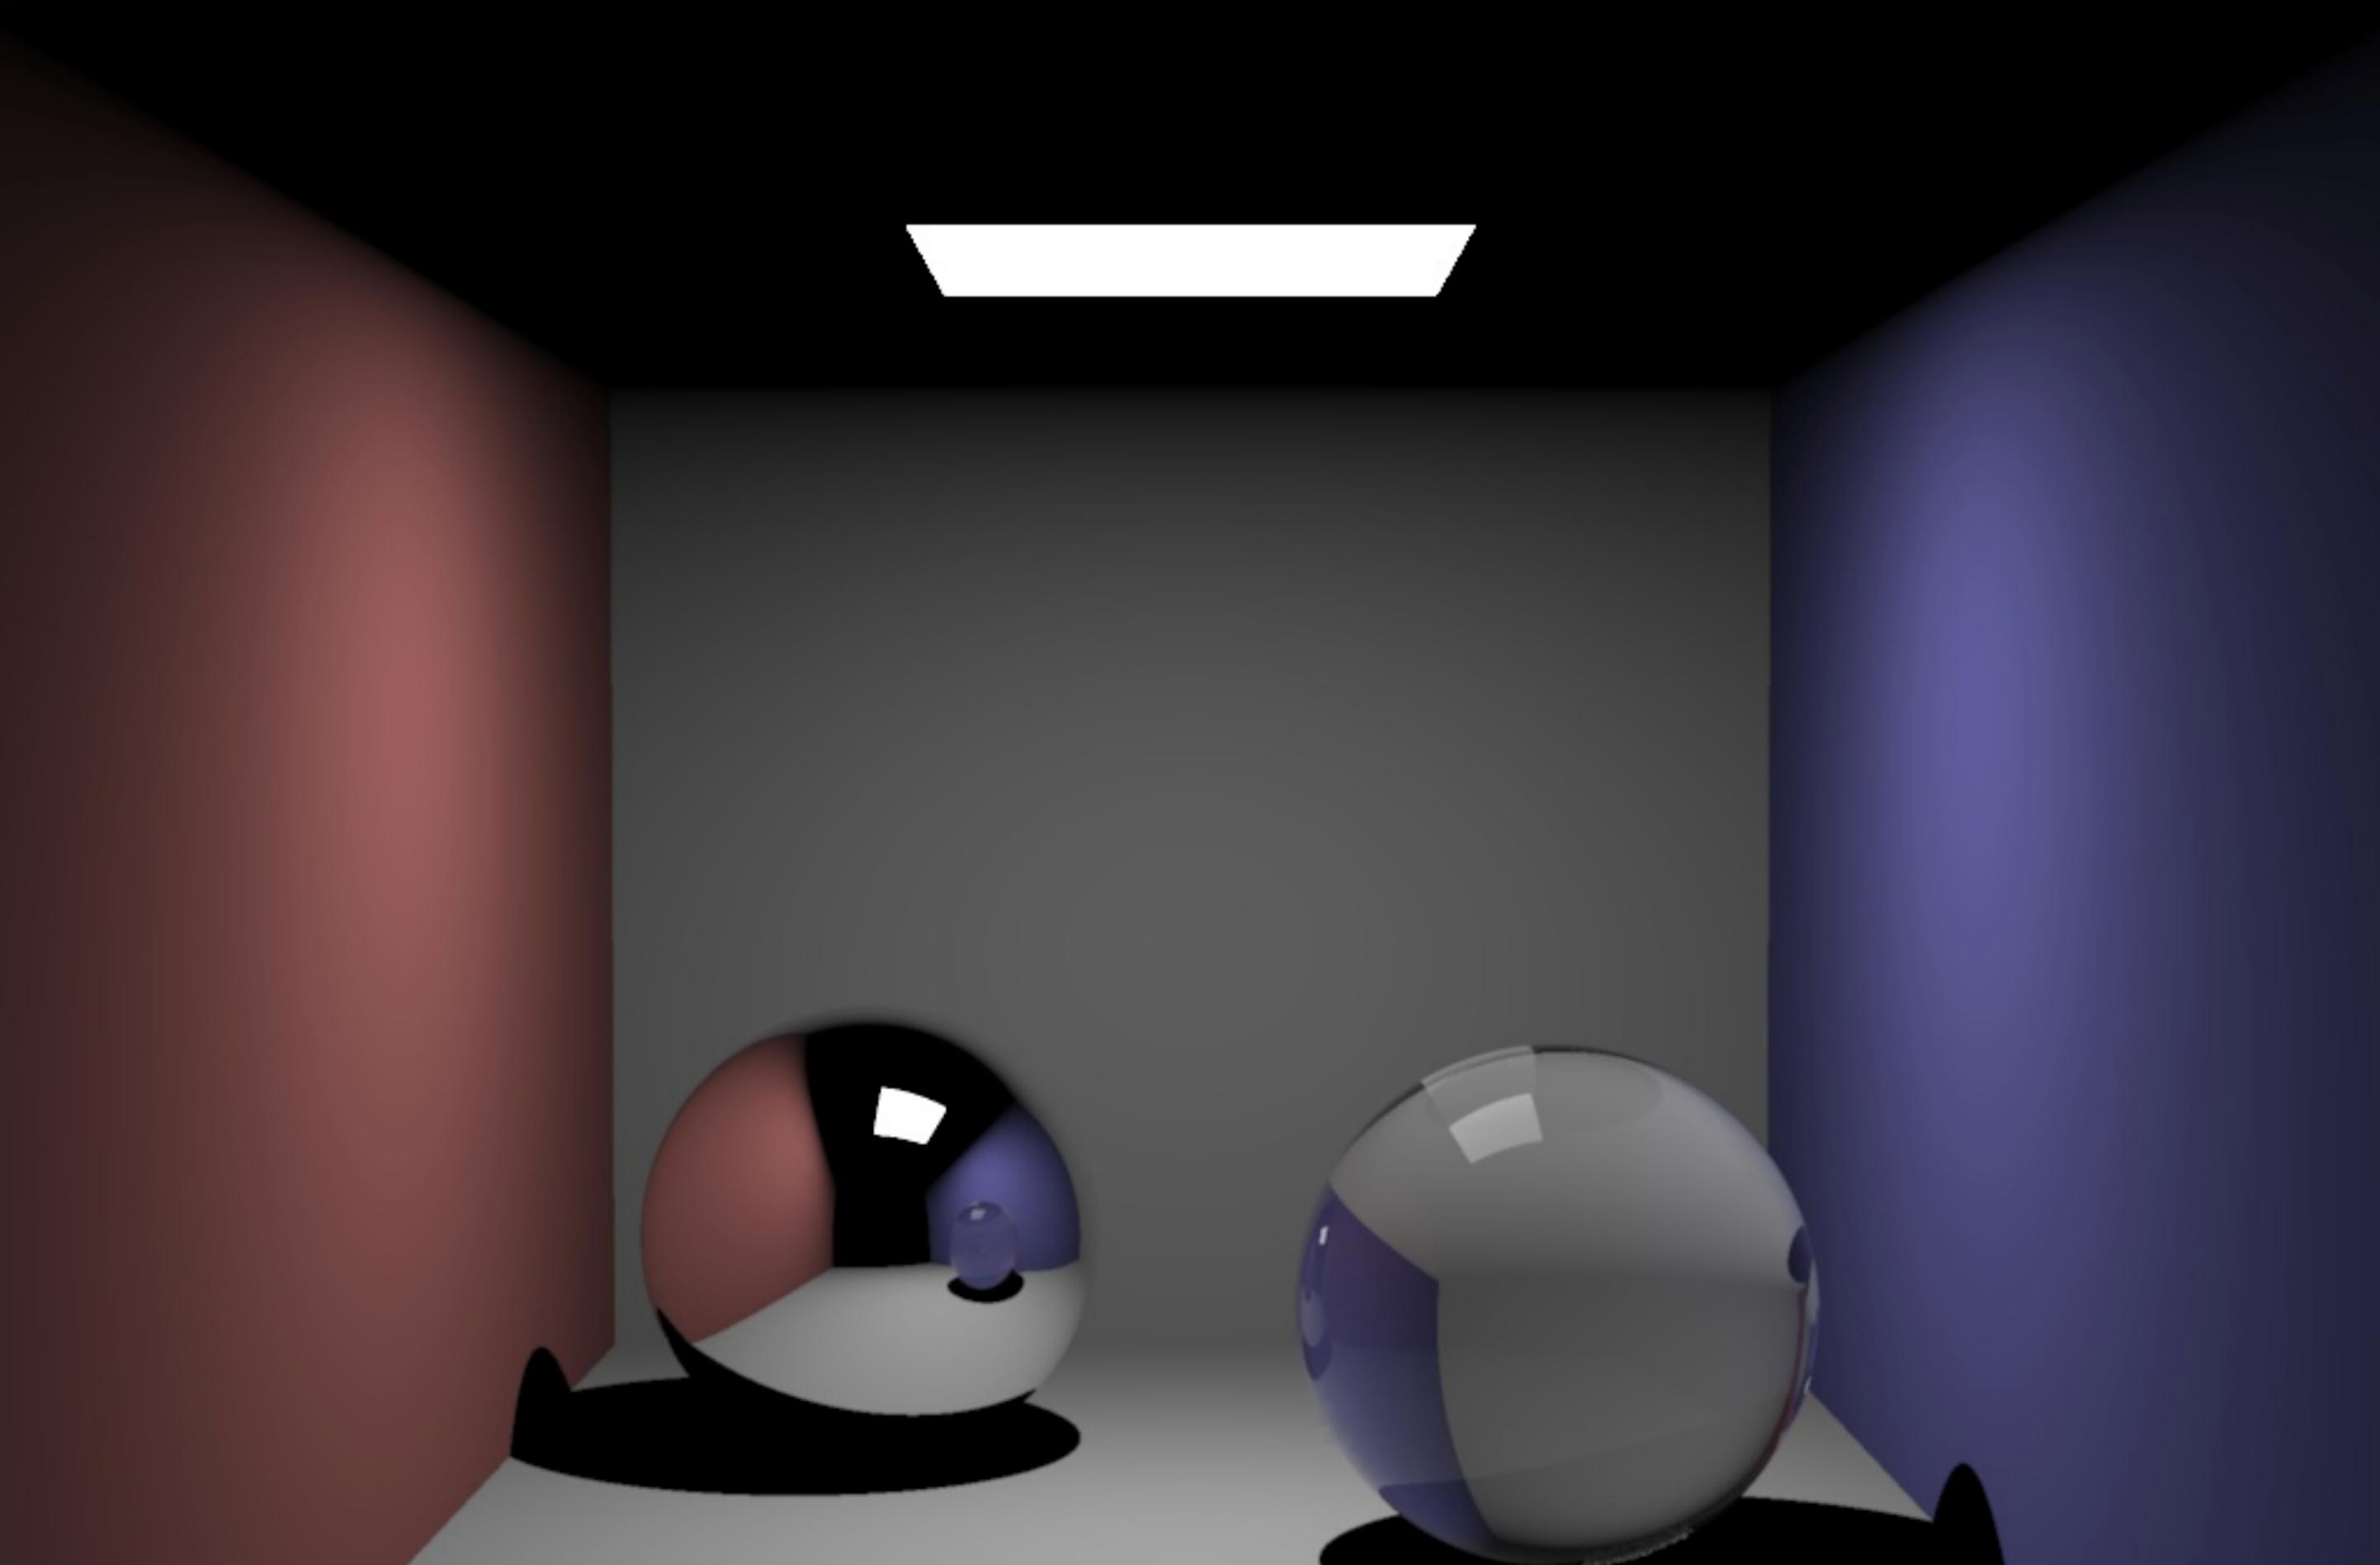
**What object is visible to the camera?**

**What light sources are visible from a point on a surface (Is a surface in shadow?)**

**What reflection is visible on a surface?**

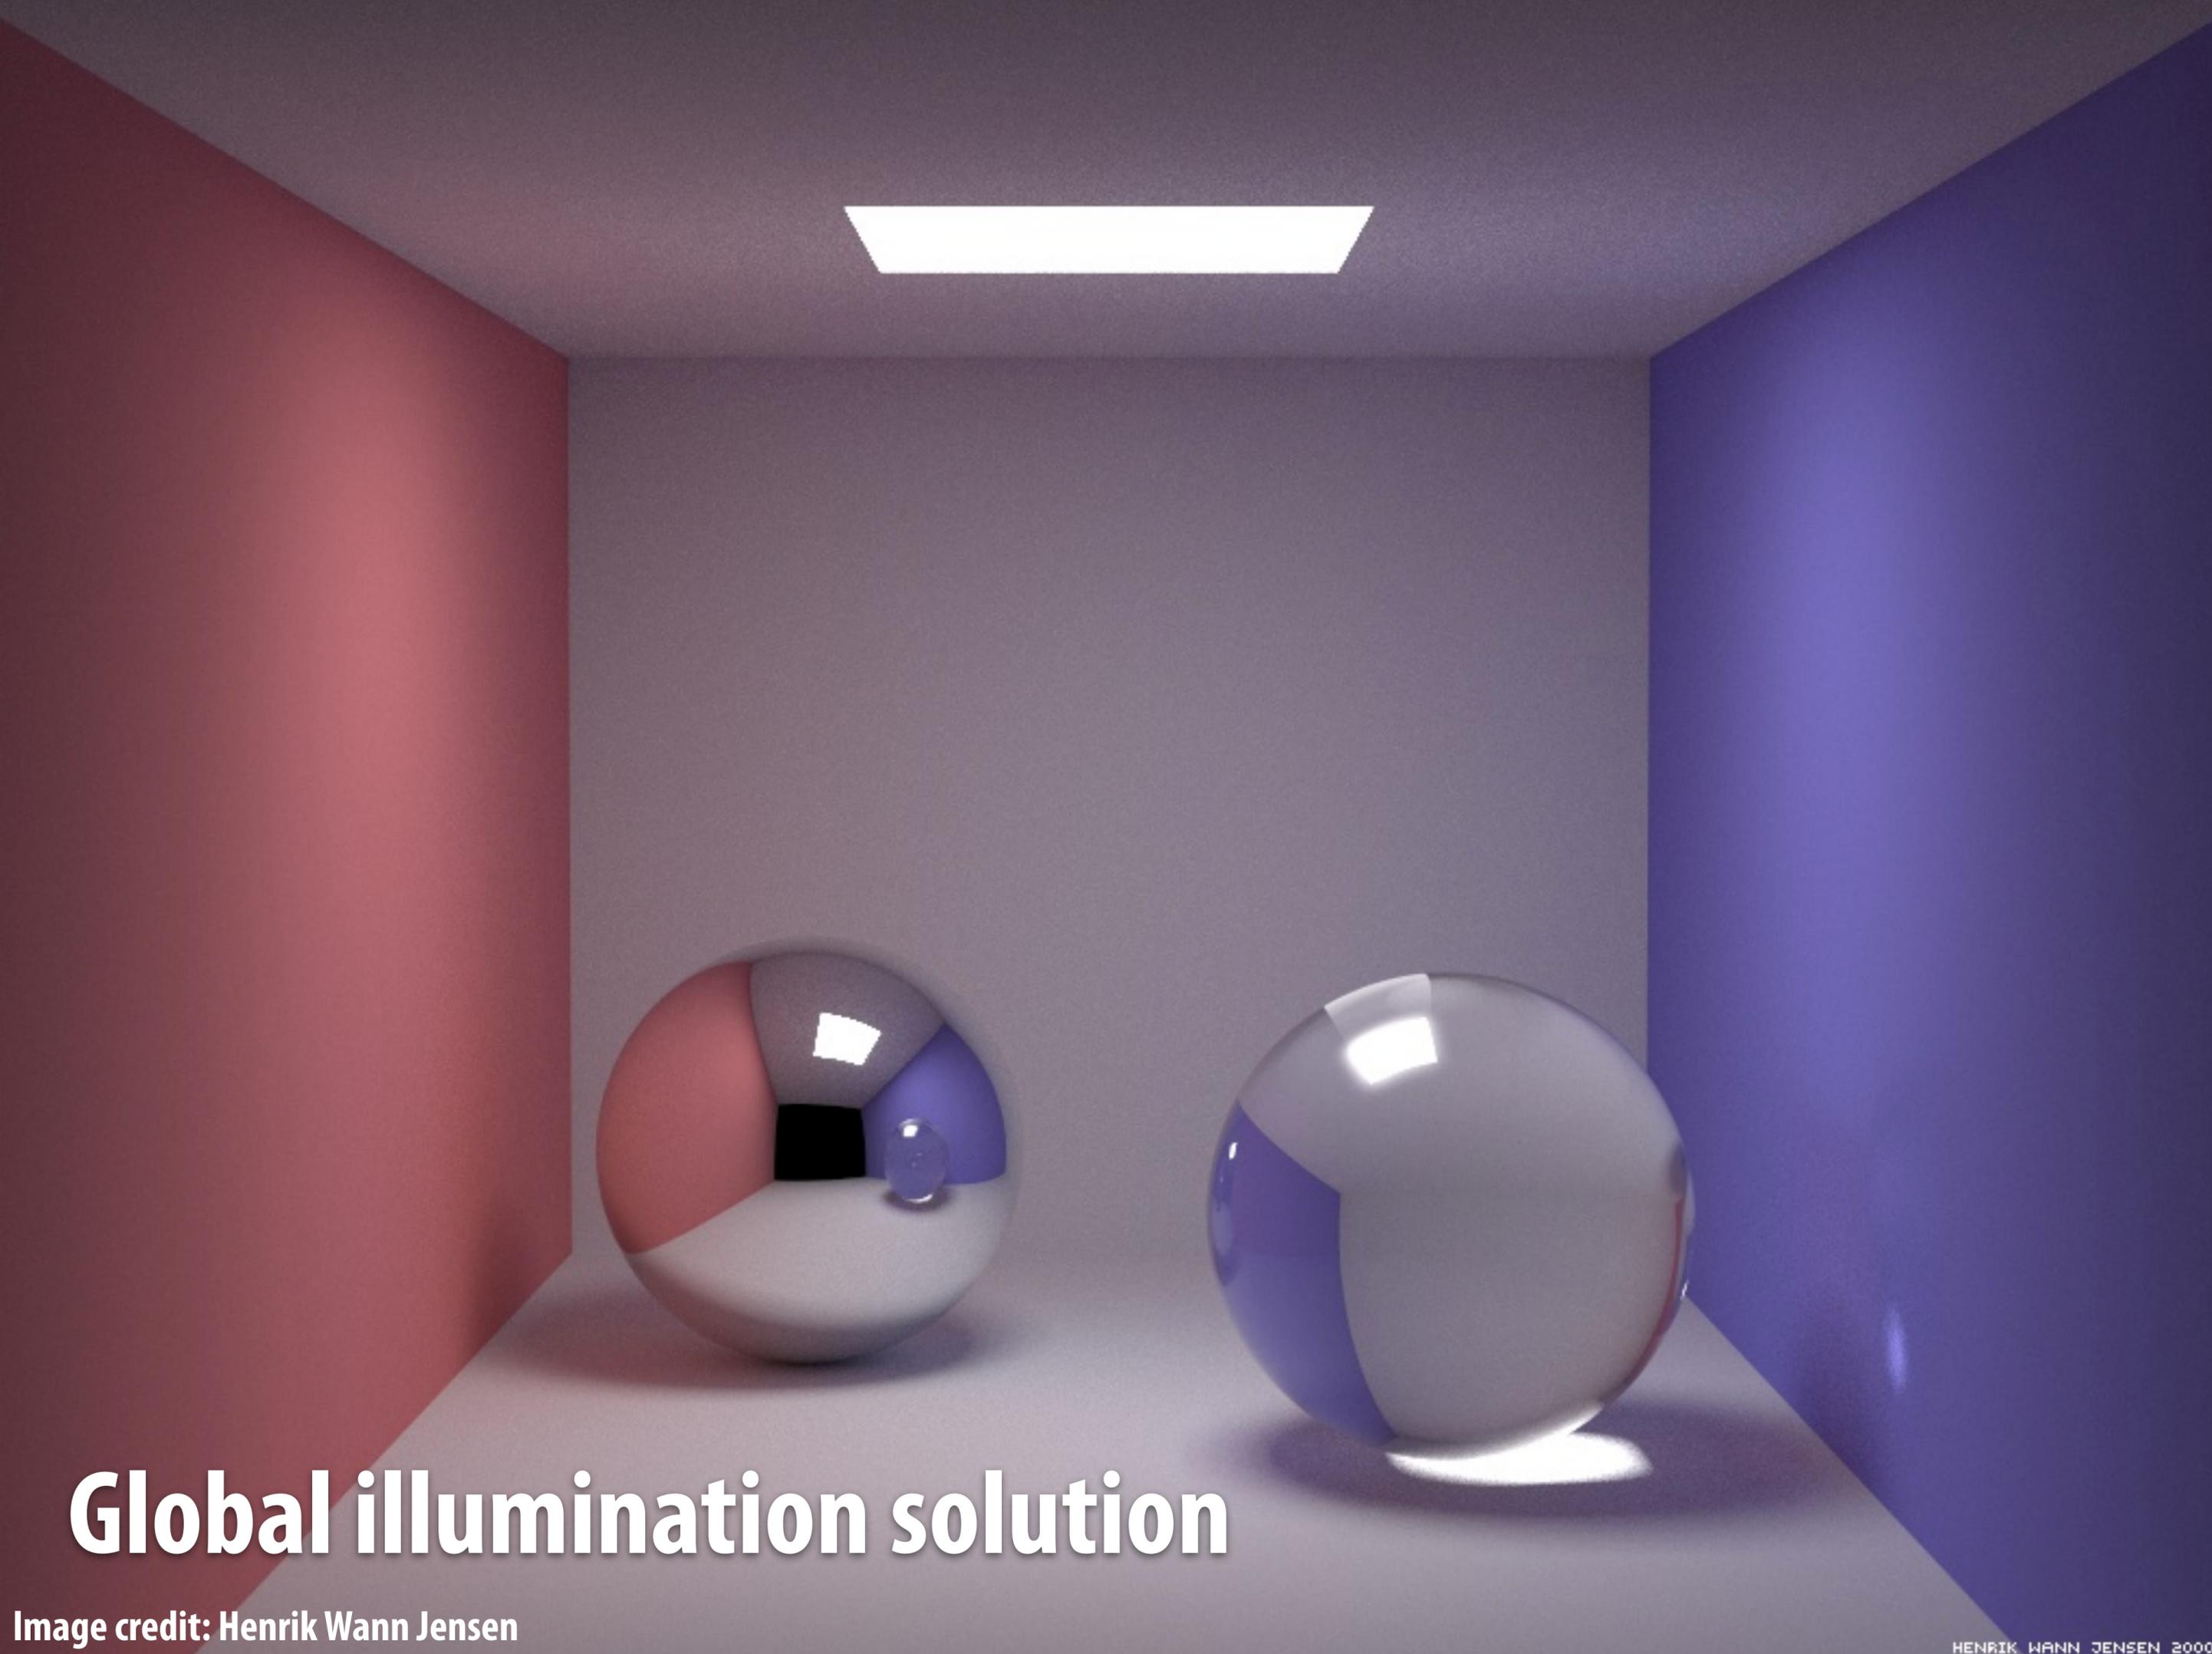


**In contrast, rasterization is a highly-specialized solution for computing visibility for a set of uniformly distributed rays originating from the same point (most often: the camera)**



# Direct illumination + reflection + transparency

Image credit: Henrik Wann Jensen



# Global illumination solution

Image credit: Henrik Wann Jensen

A photograph of a courtyard with a series of arches and columns. The scene is brightly lit, with strong shadows cast on the ground, indicating direct illumination. The architecture features a series of arches supported by columns, creating a covered walkway. The courtyard is paved with stone tiles and contains several tables and chairs, suggesting an outdoor dining area. There are plants and trees in the courtyard, and a balcony with a railing is visible in the background. The overall atmosphere is bright and sunny.

**Direct illumination**

• *p*



• p

**Sixteen-bounce global illumination**

# Increasing interest in high performance implementations of real-time ray tracing

Microsoft's DirectX Ray Tracing support / NVIDIA's DXR announced in April 2018

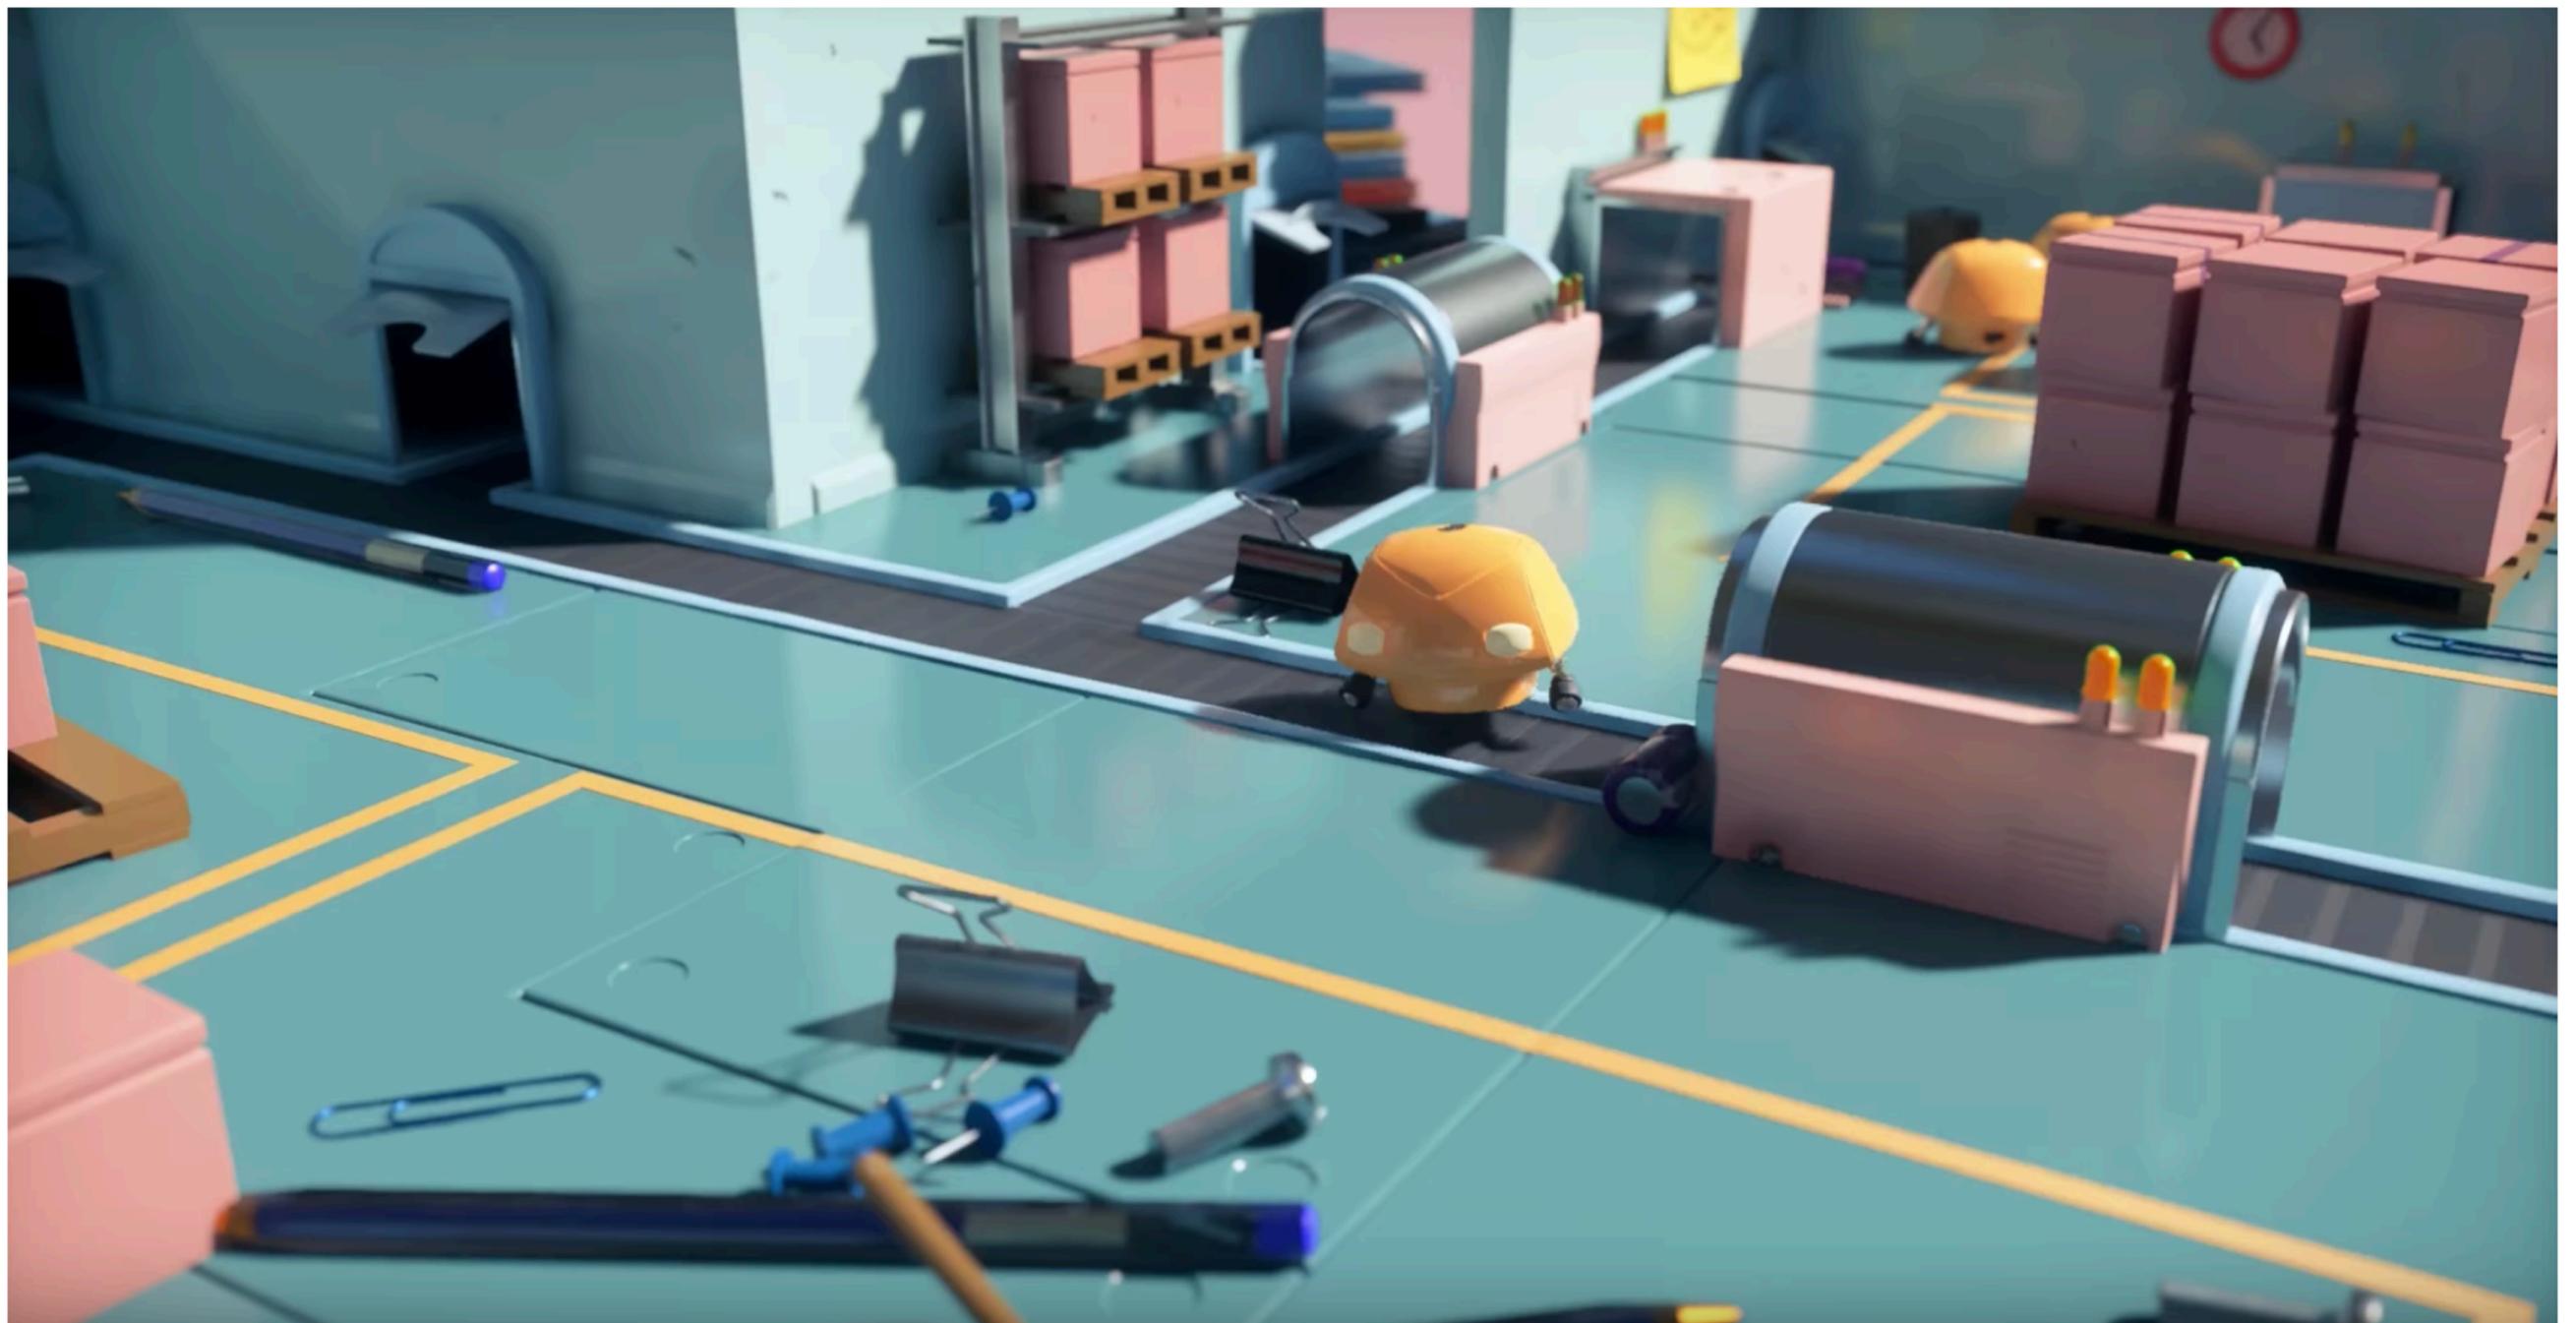


Image credit: Electronic Arts (Project PICA)

# Acknowledgements

- **Thanks to Keenan Crane, Ren Ng, and Matt Pharr for presentation resources**