

The Iajolla renderer

UCSD CSE 272

Advanced Image Synthesis

Tzu-Mao Li

From smallpt to lajolla

BachiLi / lajolla_public Public

Unwatch 1 Fork 0

Code Issues Pull requests Actions Projects Wiki Security Insights Settings

main 1 branch 0 tags

Go to file Add file Code

About

UCSD CSE 272 renderer

Readme MIT License 7 stars 1 watching 0 forks

Releases

No releases published Create a new release

Packages

No packages published Publish your first package

Contributors 2

BachiLi Tzu-Mao Li AlstonXiao Yan Xiao





BachiLi Set up volumetric path tracing templates 59c114b 19 hours ago 142 commits

cmake	Add Embree dependency	22 days ago
embree	Add Embree dependency	22 days ago
handouts	HW1 handout typo fix	5 days ago
scenes	Set up volumetric path tracing templates	19 hours ago
src	Set up volumetric path tracing templates	19 hours ago
.gitignore	Add UV & shading frame computation. Fixed sponza rendering.	16 days ago
CMakeLists.txt	Set up volumetric path tracing templates	19 hours ago
LICENSE	Initial commit	2 months ago
README.md	update readme	11 days ago

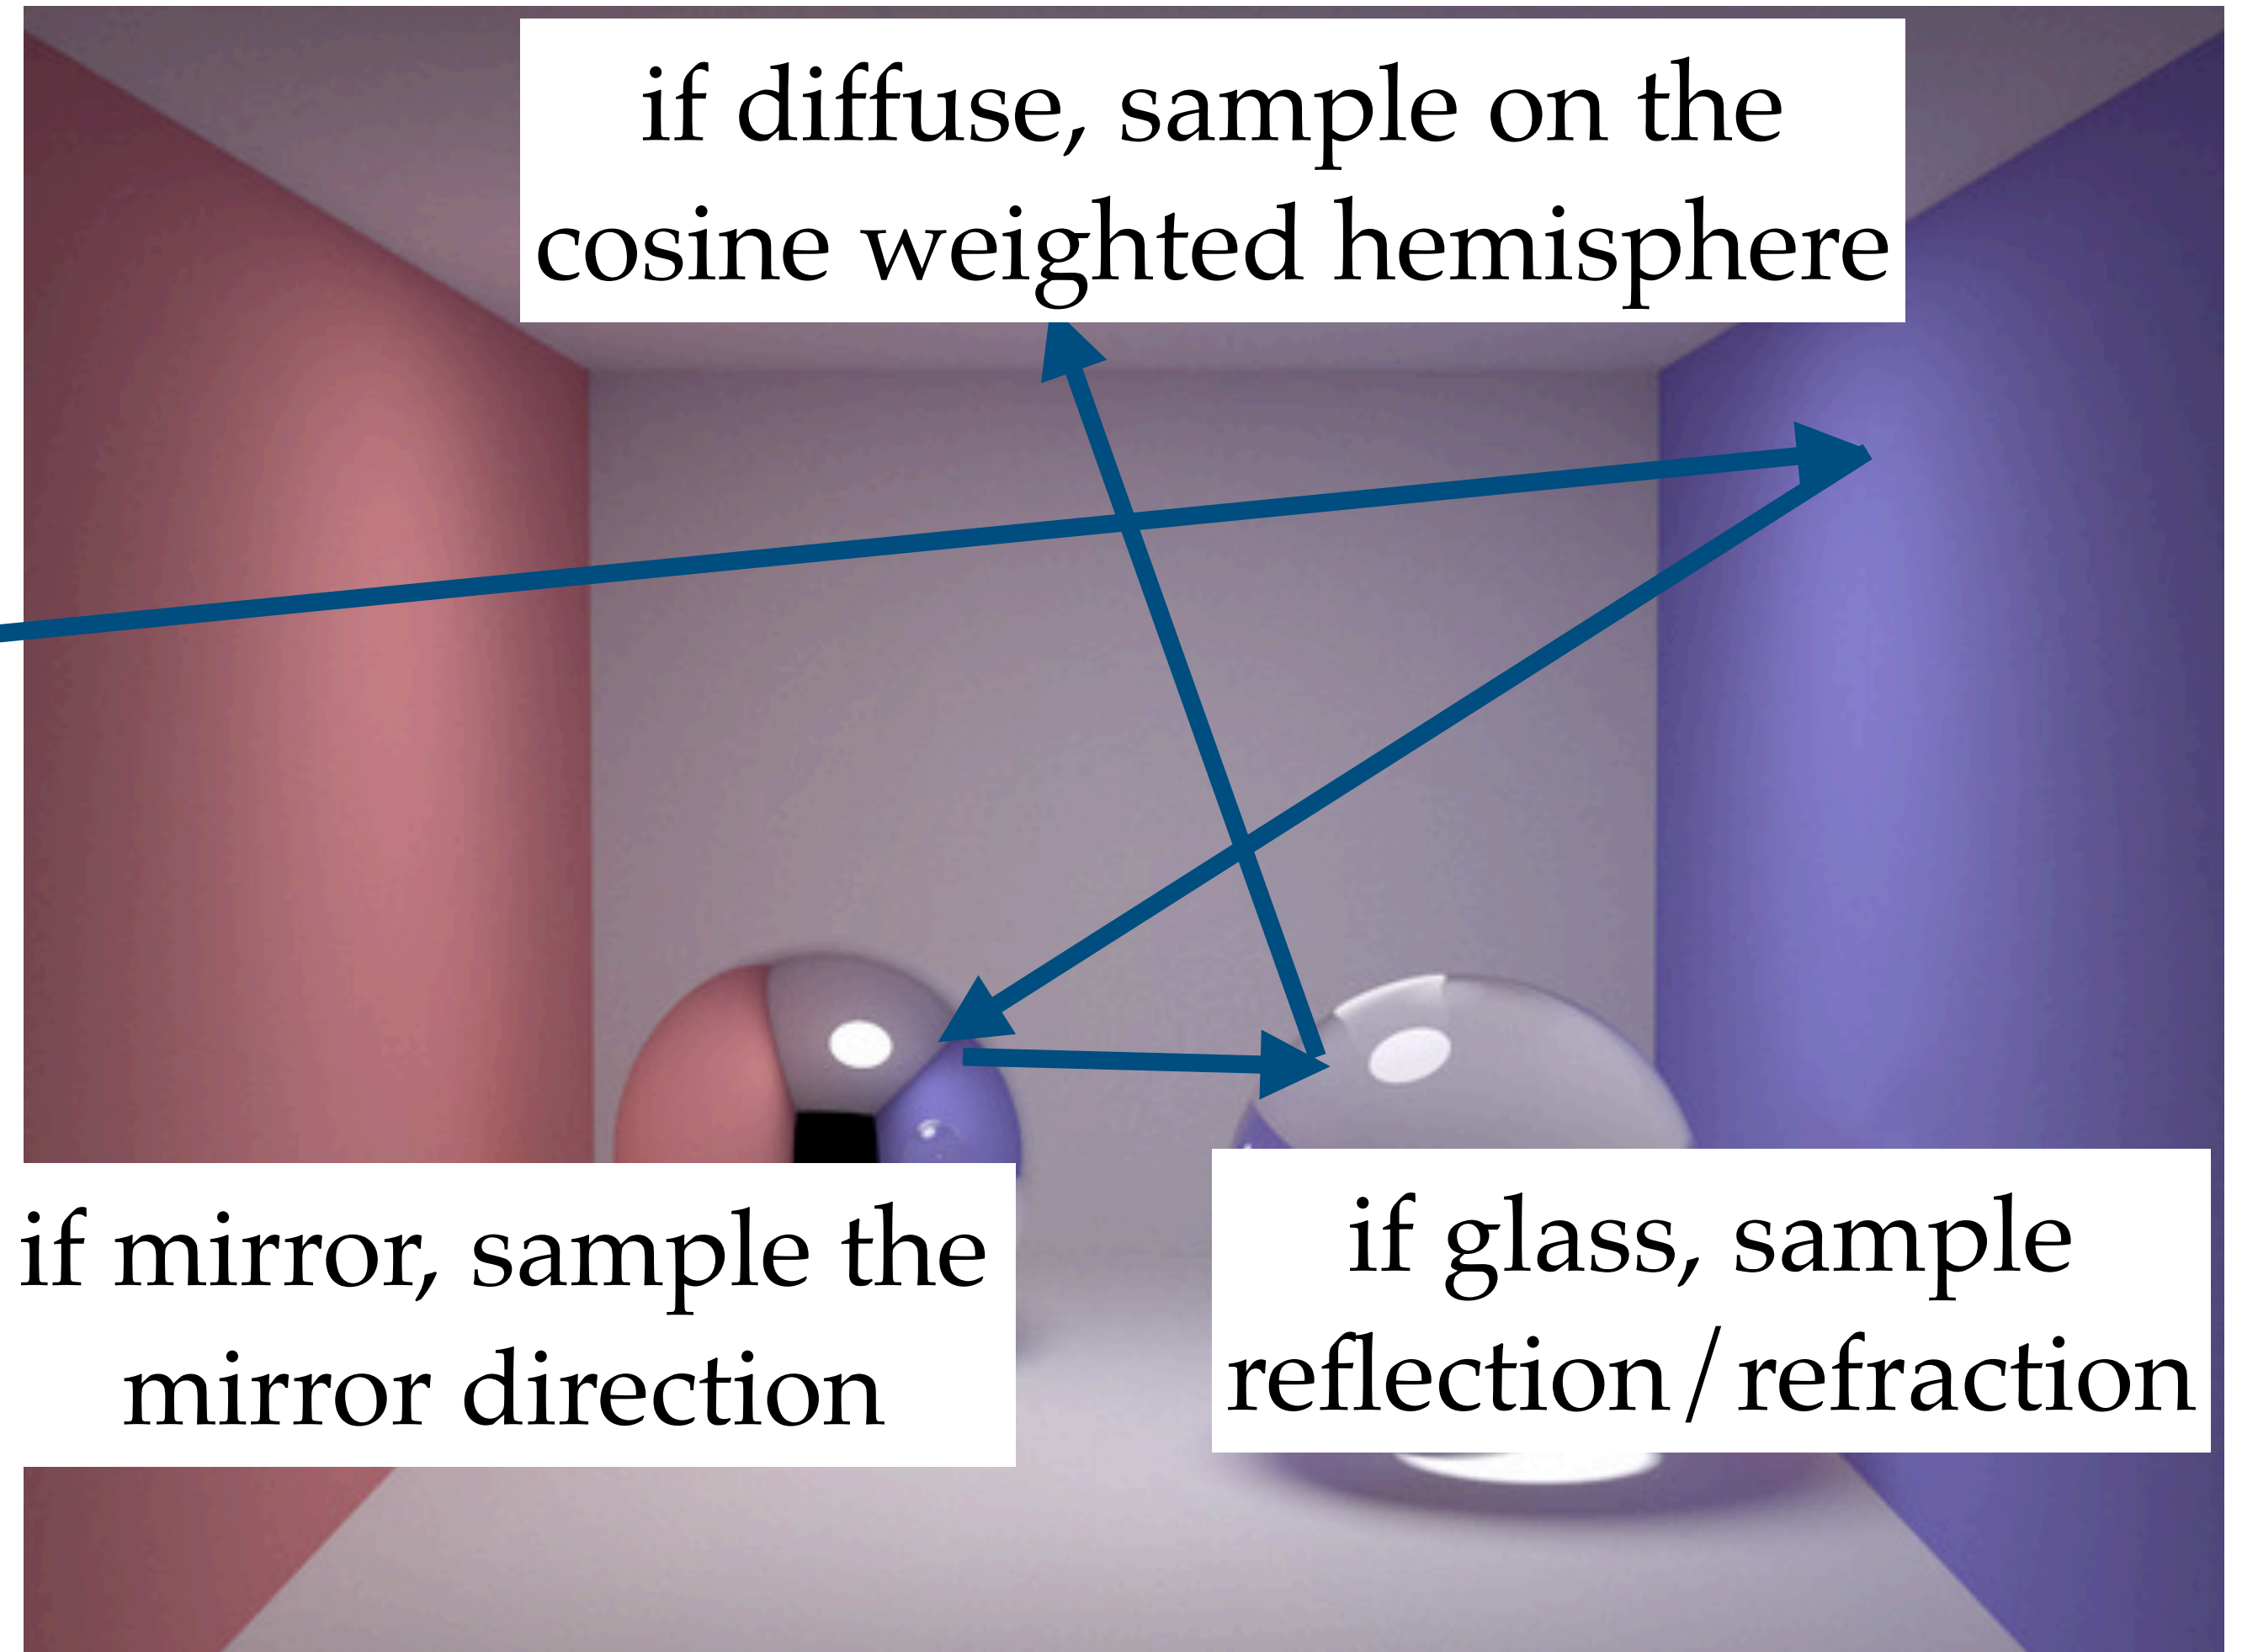

README.md

lajolla

UCSD CSE 272 renderer




Smallpt: hope to hit light with directional sampling



Next event estimation


- in addition to cosine-weighted hemisphere sampling, also sample a point on light



Next event estimation is also a change of variable

focus on the rays that hit the light source

$$\underbrace{\iint L'_e(\omega') |\omega' \cdot n| d\omega'}_{\text{red arc}}$$



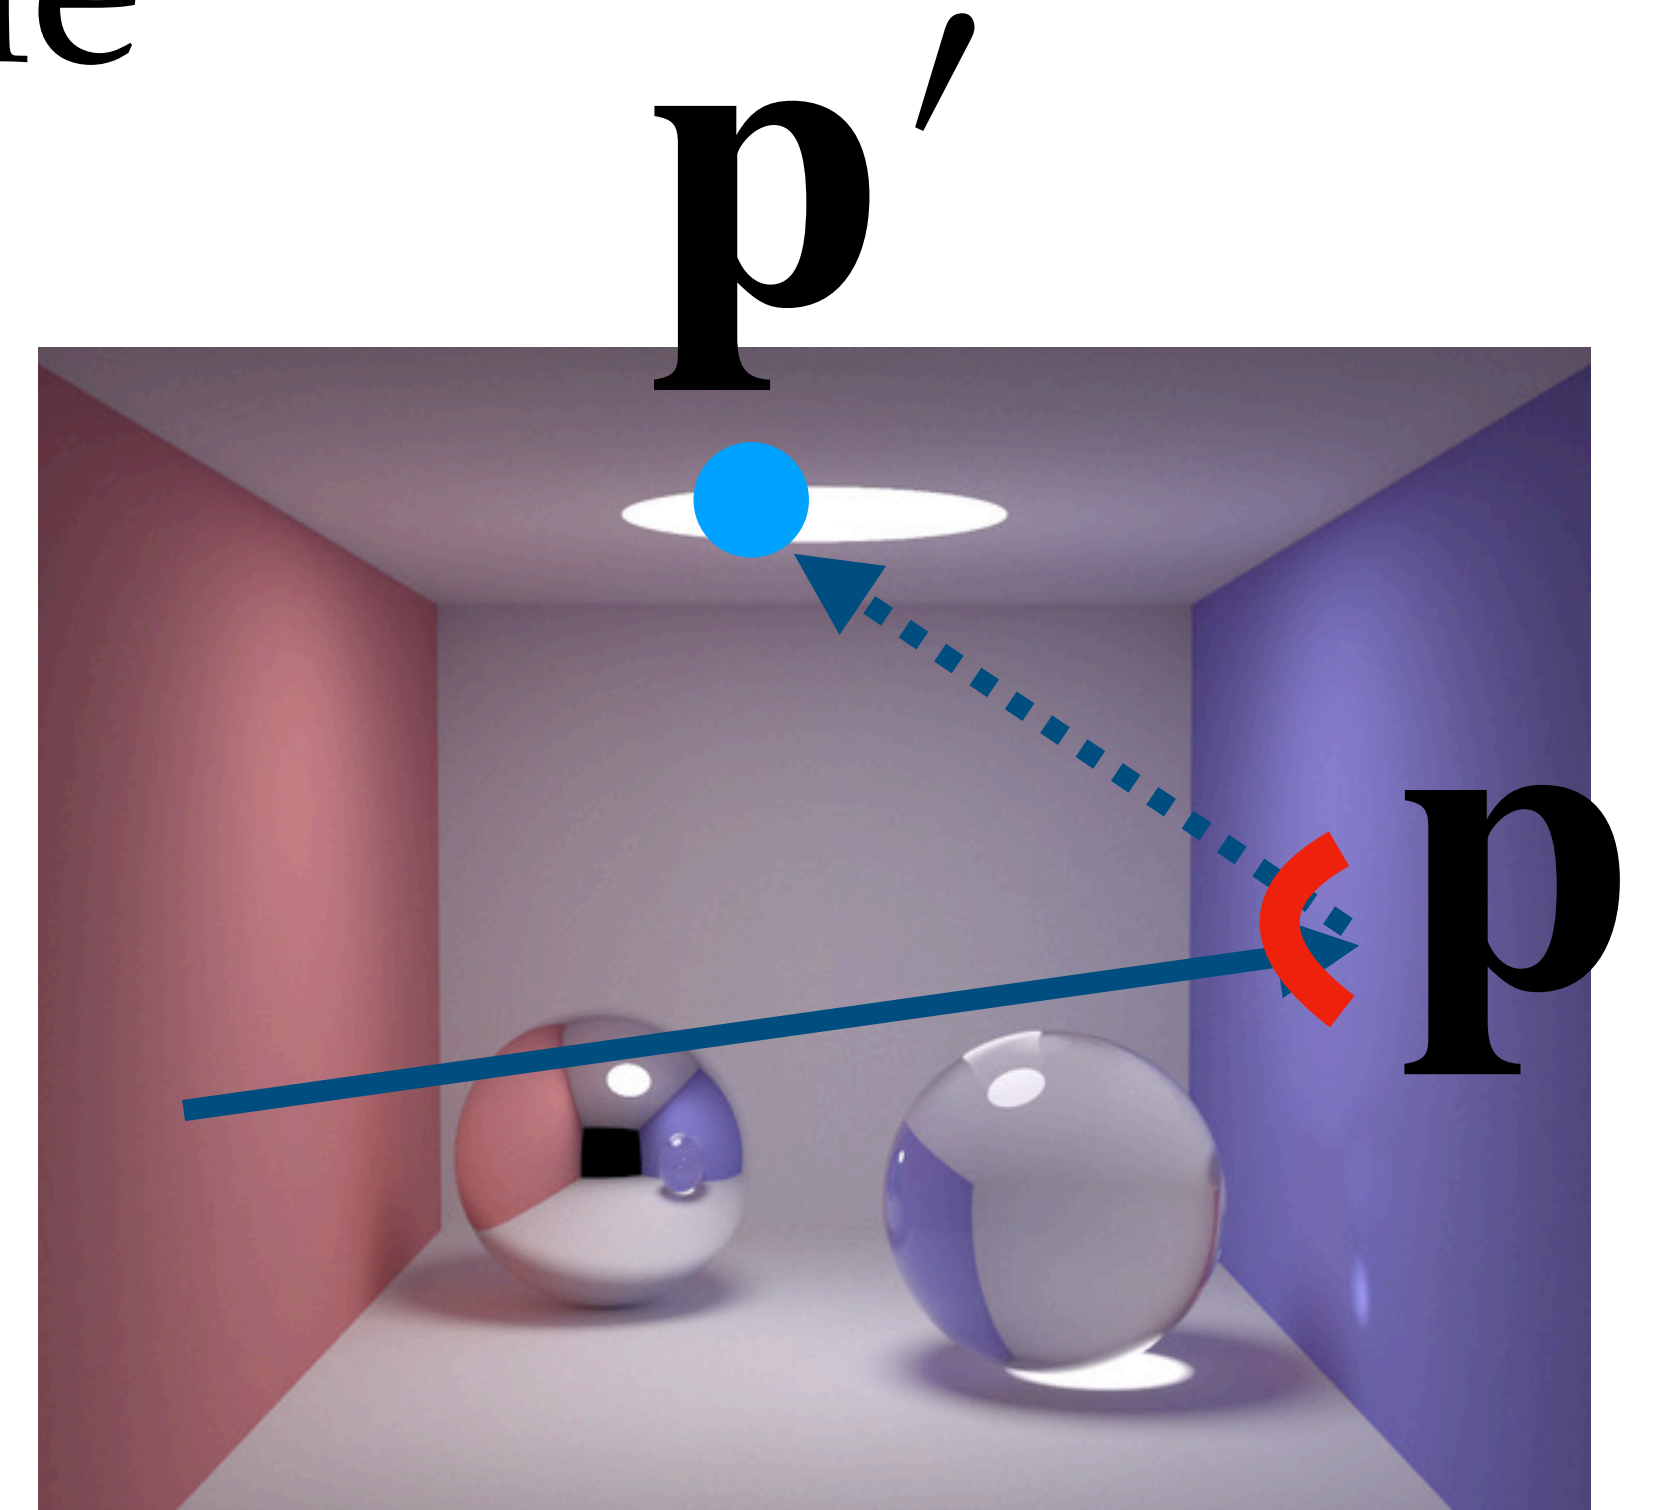

Next event estimation is also a
change of variable

focus on the rays that hit the light source

$$\iint L'_e(\omega') |\omega' \cdot n| d\omega'$$




$$\iint L'_e(\omega'(\mathbf{p}')) |\omega'(\mathbf{p}') \cdot n_{\mathbf{p}}| \frac{|\omega'(\mathbf{p}') \cdot n_{\mathbf{p}'}|}{\|\mathbf{p} - \mathbf{p}'\|^2} \text{visible}(\mathbf{p}, \mathbf{p}') d\mathbf{p}'$$




the Jacobian (often called “geometry term”)


Handling multiple bounces



Handling multiple bounces




Handling multiple bounces




Triangle-mesh light sampling:

1. pick a light based on their intensities
2. pick a triangle based on its area
3. pick a point on the triangle




see https://cseweb.ucsd.edu/~tzli/cse272/wi2023/lectures/triangle_sampling.pdf
for notes on triangle sampling

Next event estimation is good at small lights




next event estimation
(64 samples per pixel)




cosine weighted hemisphere
(64 samples per pixel)

When the point is close to the light,
cosine-weighted hemisphere sampling is better



next event estimation




cosine-weighted
hemisphere sampling

Multiple importance sampling: combining next event estimation and hemisphere sampling

$$\begin{aligned} & \iint L'(\omega') |\omega' \cdot n| d\omega' \\ & \approx \frac{1}{N} \sum_{i=1}^N \frac{L'(\omega'_i) |\omega'_i \cdot n|}{p_{\text{nee}}(\omega'_i)} \\ & \approx \frac{1}{N} \sum_{j=1}^N \frac{L'(\omega'_j) |\omega'_j \cdot n|}{p_{\text{hemi}}(\omega'_j)} \end{aligned}$$

?

?



Multiple importance sampling: combining next event estimation and hemisphere sampling

- idea: assign higher weight when p is high

$$\frac{1}{N} \left(\sum_{i=1}^N w_i^n \frac{L'(\omega'_i) |\omega'_i \cdot n|}{p_{nee}(\omega'_i)} + \sum_{j=1}^N w_j^h \frac{L'(\omega'_j) |\omega'_j \cdot n|}{p_{hemi}(\omega'_j)} \right)$$

Multiple importance sampling: combining next event estimation and hemisphere sampling

- idea: assign higher weight when p is high

$$\frac{1}{N} \left(\sum_{i=1}^N w_i^n \frac{L'(\omega'_i) |\omega'_i \cdot n|}{p_{\text{nee}}(\omega'_i)} + \sum_{j=1}^N w_j^h \frac{L'(\omega'_j) |\omega'_j \cdot n|}{p_{\text{hemi}}(\omega'_j)} \right)$$

$$w_i^n = \frac{p_{\text{nee}}(\omega_i)}{p_{\text{nee}}(\omega_i) + p_{\text{hemi}}(\omega_i)}$$

Multiple importance sampling: combining next event estimation and hemisphere sampling

- idea: assign higher weight when p is high

$$\frac{1}{N} \left(\sum_{i=1}^N w_i^n \frac{L'(\omega'_i) |\omega'_i \cdot n|}{p_{\text{nee}}(\omega'_i)} + \sum_{j=1}^N w_j^h \frac{L'(\omega'_j) |\omega'_j \cdot n|}{p_{\text{hemi}}(\omega'_j)} \right)$$

$$w_i^n = \frac{p_{\text{nee}}(\omega_i)}{p_{\text{nee}}(\omega_i) + p_{\text{hemi}}(\omega_i)}$$

$$w_j^h = \frac{p_{\text{hemi}}(\omega_j)}{p_{\text{nee}}(\omega_j) + p_{\text{hemi}}(\omega_j)}$$

Multiple importance sampling: combining next event estimation and hemisphere sampling

- idea: assign higher weight when p is high

$$\frac{1}{N} \left(\sum_{i=1}^N w_i^n \frac{L'(\omega'_i) |\omega'_i \cdot n|}{p_{\text{nee}}(\omega'_i)} + \sum_{j=1}^N w_j^h \frac{L'(\omega'_j) |\omega'_j \cdot n|}{p_{\text{hemi}}(\omega'_j)} \right)$$


$$w_i^n = \frac{p_{\text{nee}}(\omega_i)}{p_{\text{nee}}(\omega_i) + p_{\text{hemi}}(\omega_i)}$$

$$w_j^h = \frac{p_{\text{hemi}}(\omega_j)}{p_{\text{nee}}(\omega_j) + p_{\text{hemi}}(\omega_j)}$$


see https://github.com/BachiLi/lajolla_public/blob/main/src/path_tracing.h for the implementation

https://cseweb.ucsd.edu/~tzli/cse168/lectures/12_multiple_importance_sampling.pdf for the math


MIS combines the best of both worlds



next event estimation




cosine-weighted
hemisphere sampling



multiple importance sampling

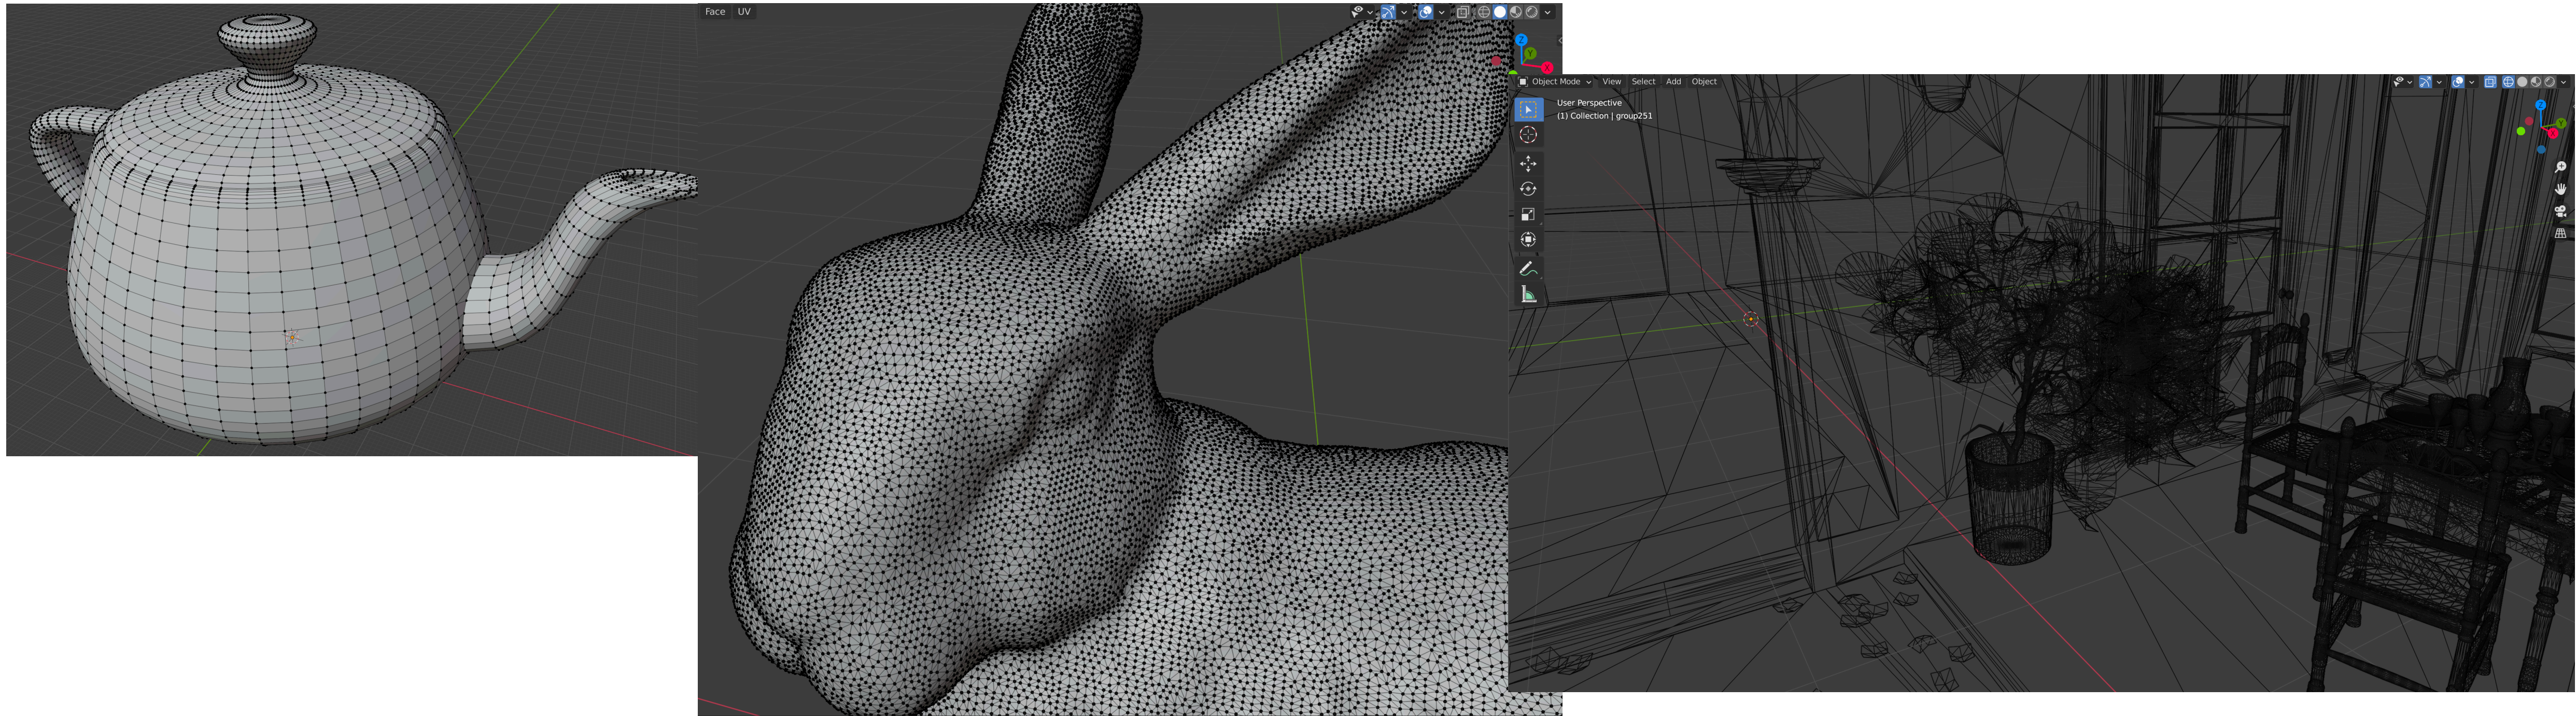
Smallpt: sphere geometry only

```
1. Sphere spheres[] = {//Scene: radius, position, emission, color, material
2.   Sphere(1e5, Vec( 1e5+1, 40.8, 81.6), Vec(), Vec(.75, .25, .25), DIFF), //Left
3.   Sphere(1e5, Vec(-1e5+99, 40.8, 81.6), Vec(), Vec(.25, .25, .75), DIFF), //Right
4.   Sphere(1e5, Vec(50, 40.8, 1e5), Vec(), Vec(.75, .75, .75), DIFF), //Back
5.   Sphere(1e5, Vec(50, 40.8, -1e5+170), Vec(), Vec(), DIFF), //Frnt
6.   Sphere(1e5, Vec(50, 1e5, 81.6), Vec(), Vec(.75, .75, .75), DIFF), //Botm
7.   Sphere(1e5, Vec(50, -1e5+81.6, 81.6), Vec(), Vec(.75, .75, .75), DIFF), //Top
8.   Sphere(16.5, Vec(27, 16.5, 47), Vec(), Vec(1, 1, 1)*.999, SPEC), //Mirr
9.   Sphere(16.5, Vec(73, 16.5, 78), Vec(), Vec(1, 1, 1)*.999, REFR), //Glas
10.  Sphere(600, Vec(50, 681.6-.27, 81.6), Vec(12, 12, 12), Vec(), DIFF) //Lite
11. };
```




More commonly used geometry primitive: triangle mesh

quiz: why?




Ray-triangle intersection

quiz: how do we solve for intersection?



$$\mathbf{x} = \mathbf{o} + t \cdot \mathbf{d}$$

P_0



P_2

P_1

$$\mathbf{x} = (1 - b_1 - b_2)\mathbf{P}_0 + b_1\mathbf{P}_1 + b_2\mathbf{P}_2$$

Ray-triangle intersection

$$\mathbf{o} + t \cdot \mathbf{d} = (1 - b_1 - b_2)\mathbf{P}_0 + b_1\mathbf{P}_1 + b_2\mathbf{P}_2$$

$$o_x + t \cdot d_x = (1 - b_1 - b_2)P_{0_x} + b_1P_{1_x} + b_2P_{2_x}$$

$$o_y + t \cdot d_y = (1 - b_1 - b_2)P_{0_y} + b_1P_{1_y} + b_2P_{2_y}$$


$$o_z + t \cdot d_z = (1 - b_1 - b_2)P_{0_z} + b_1P_{1_z} + b_2P_{2_z}$$

3 unknowns (t, b1, b2), 3 linear equations

Lajolla supports triangle meshes and spheres

A triangle mesh representing a quad:

```
positions = {{-1, -1, 0},  
             { 1, -1, 0},  
             {-1,  1, 0},  
             { 1,  1, 0}};  
indices = {{0, 1, 2},  
           {2, 1, 3}}
```



quiz: why don't we just represent everything as individual triangles?

Lajolla supports triangle meshes and spheres

```
struct ShapeBase {  
    int material_id = -1;  
    // ...  
};
```

```
struct Sphere : public ShapeBase {  
    Vector3 position;  
    Real radius;  
};
```

```
struct TriangleMesh : public ShapeBase {  
    std::vector<Vector3> positions;  
    std::vector<Vector3i> indices;  
    std::vector<Vector3> normals;  
    std::vector<Vector2> uvs;  
    // ...  
};
```

```
using Shape = std::variant<Sphere, TriangleMesh>;
```

see HW0 for the reasoning of using std::variant

Mitsuba scene format

```
<shape type="sphere">  
    <point name="center" x="0" y="0" z="0"/>  
    <float name="radius" value="1.0"/>  
    <!-- ... -->  
</shape>
```

```
<shape type="obj">  
    <string name="filename" value="meshes/cbox_floor.obj"/>  
    <!-- ... -->  
</shape>
```

Wavefront obj file (ASCII)


https://en.wikipedia.org/wiki/Wavefront_.obj_file

Mitsuba serialized mesh (binary)


```
<shape type="serialized">  
    <string name="filename" value="matpreview.serialized"/>  
    <integer name="shapeIndex" value="2"/>  
    <transform name="toWorld">  
        <matrix value="0.614046 0.614047 0 -1.78814e-07 -0.614047 0 0 0 1"/>  
        <translate z="0.01"/>  
    </transform>  
    <!-- ... -->  
</shape>
```


Smallpt: loop over all spheres to test intersections


```
1. inline bool intersect(const Ray &r, double &t, int &id){
2.     double n=sizeof(spheres)/sizeof(Sphere), d, inf=t=1e20;
3.     for(int i=int(n);i--;) if((d=spheres[i].intersect(r))&&d<t){t=d;id=i;}
4.     return t<inf;
5. }
```



Lots of triangles:
how to go through them efficiently?




$$\mathbf{x} = \mathbf{o} + t \cdot \mathbf{d}$$



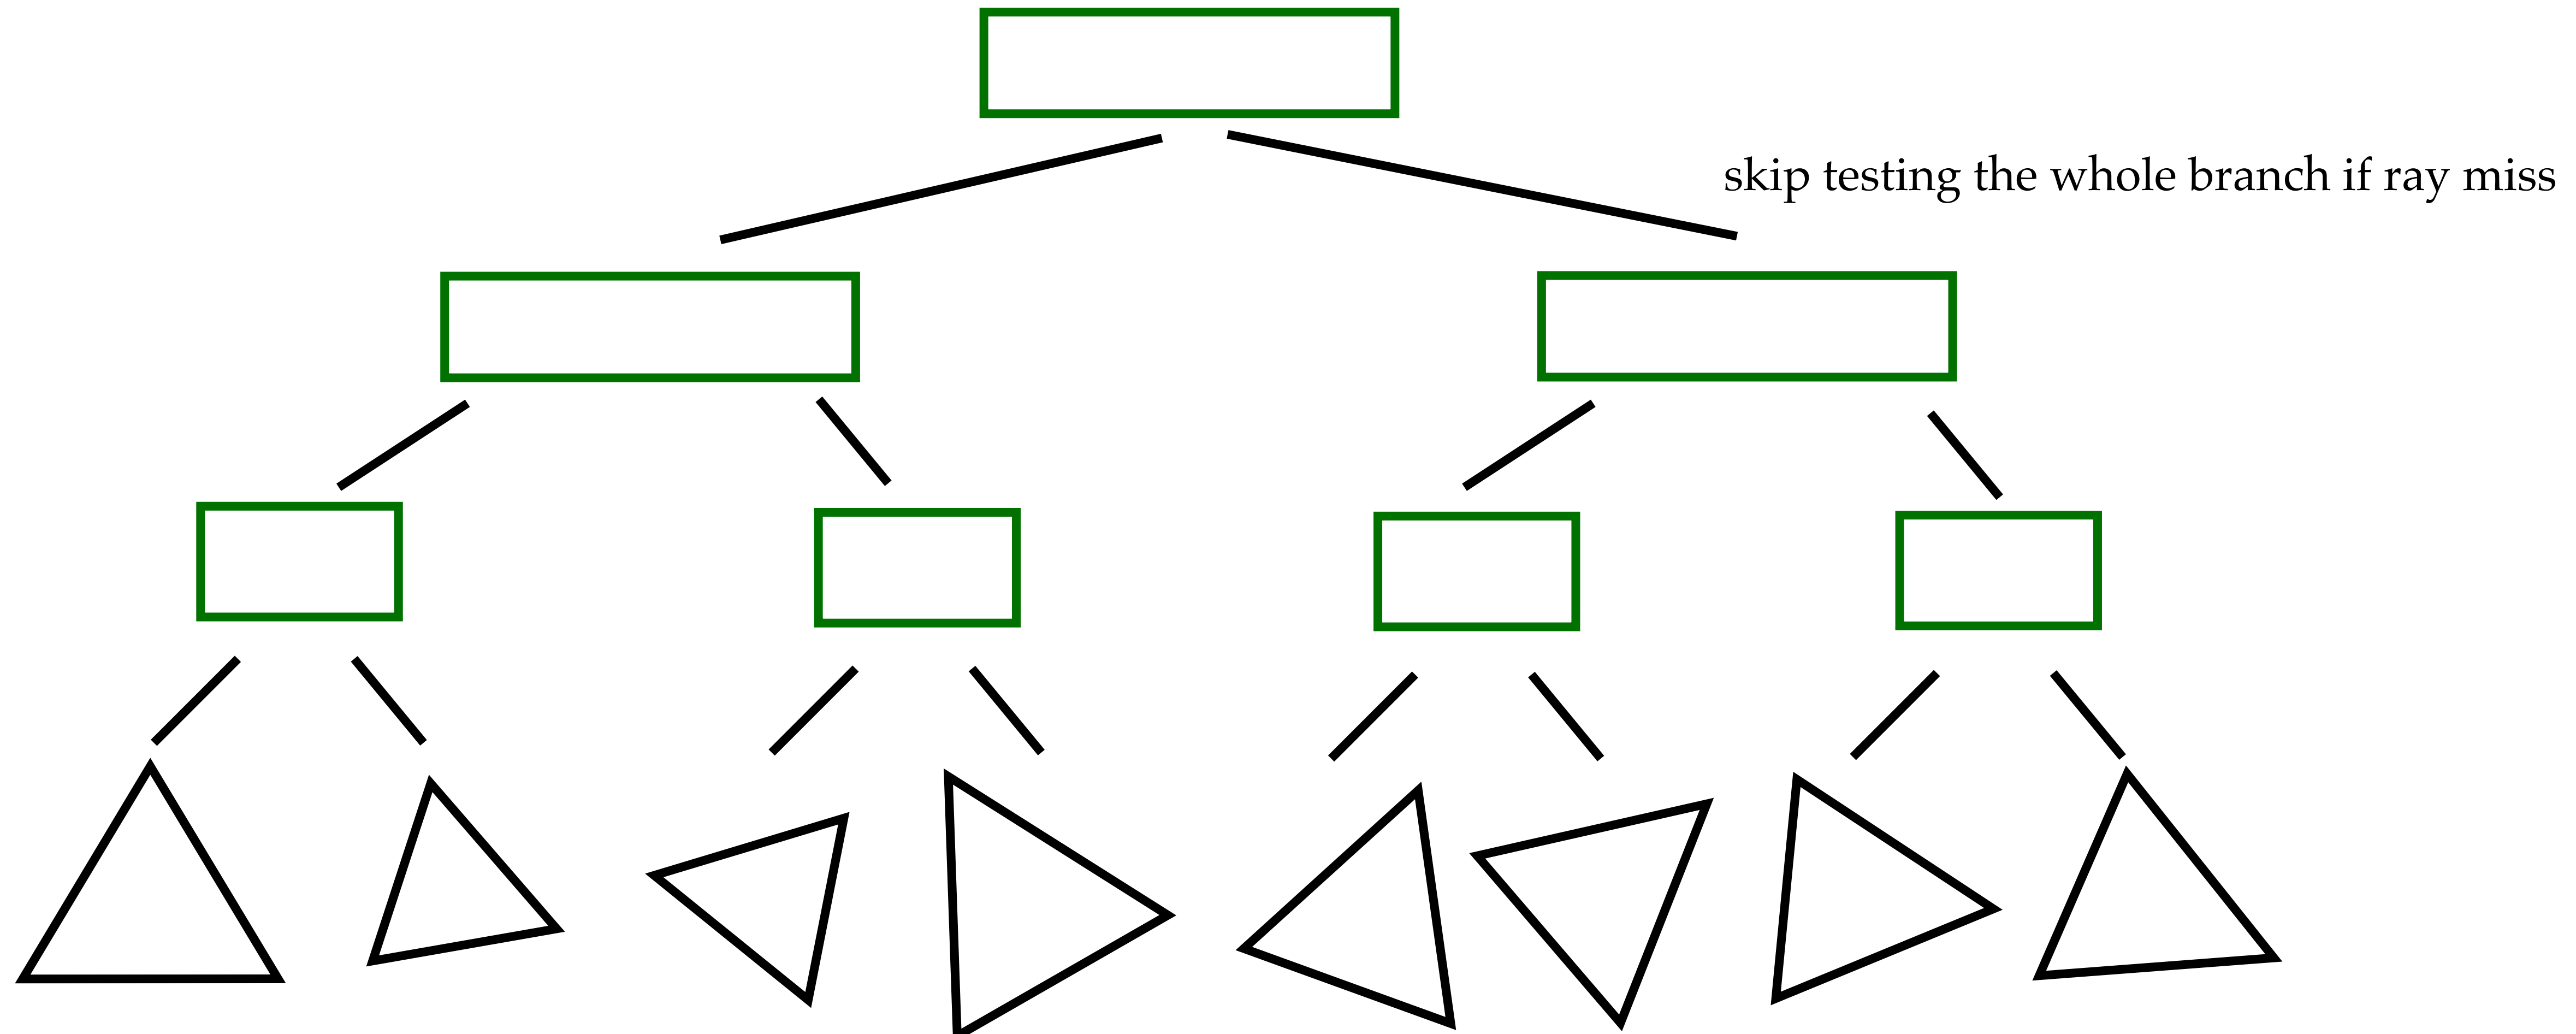
Bounding Volumes Hierarchy

- idea: test a group of triangles at a time by looking at their bounding volumes



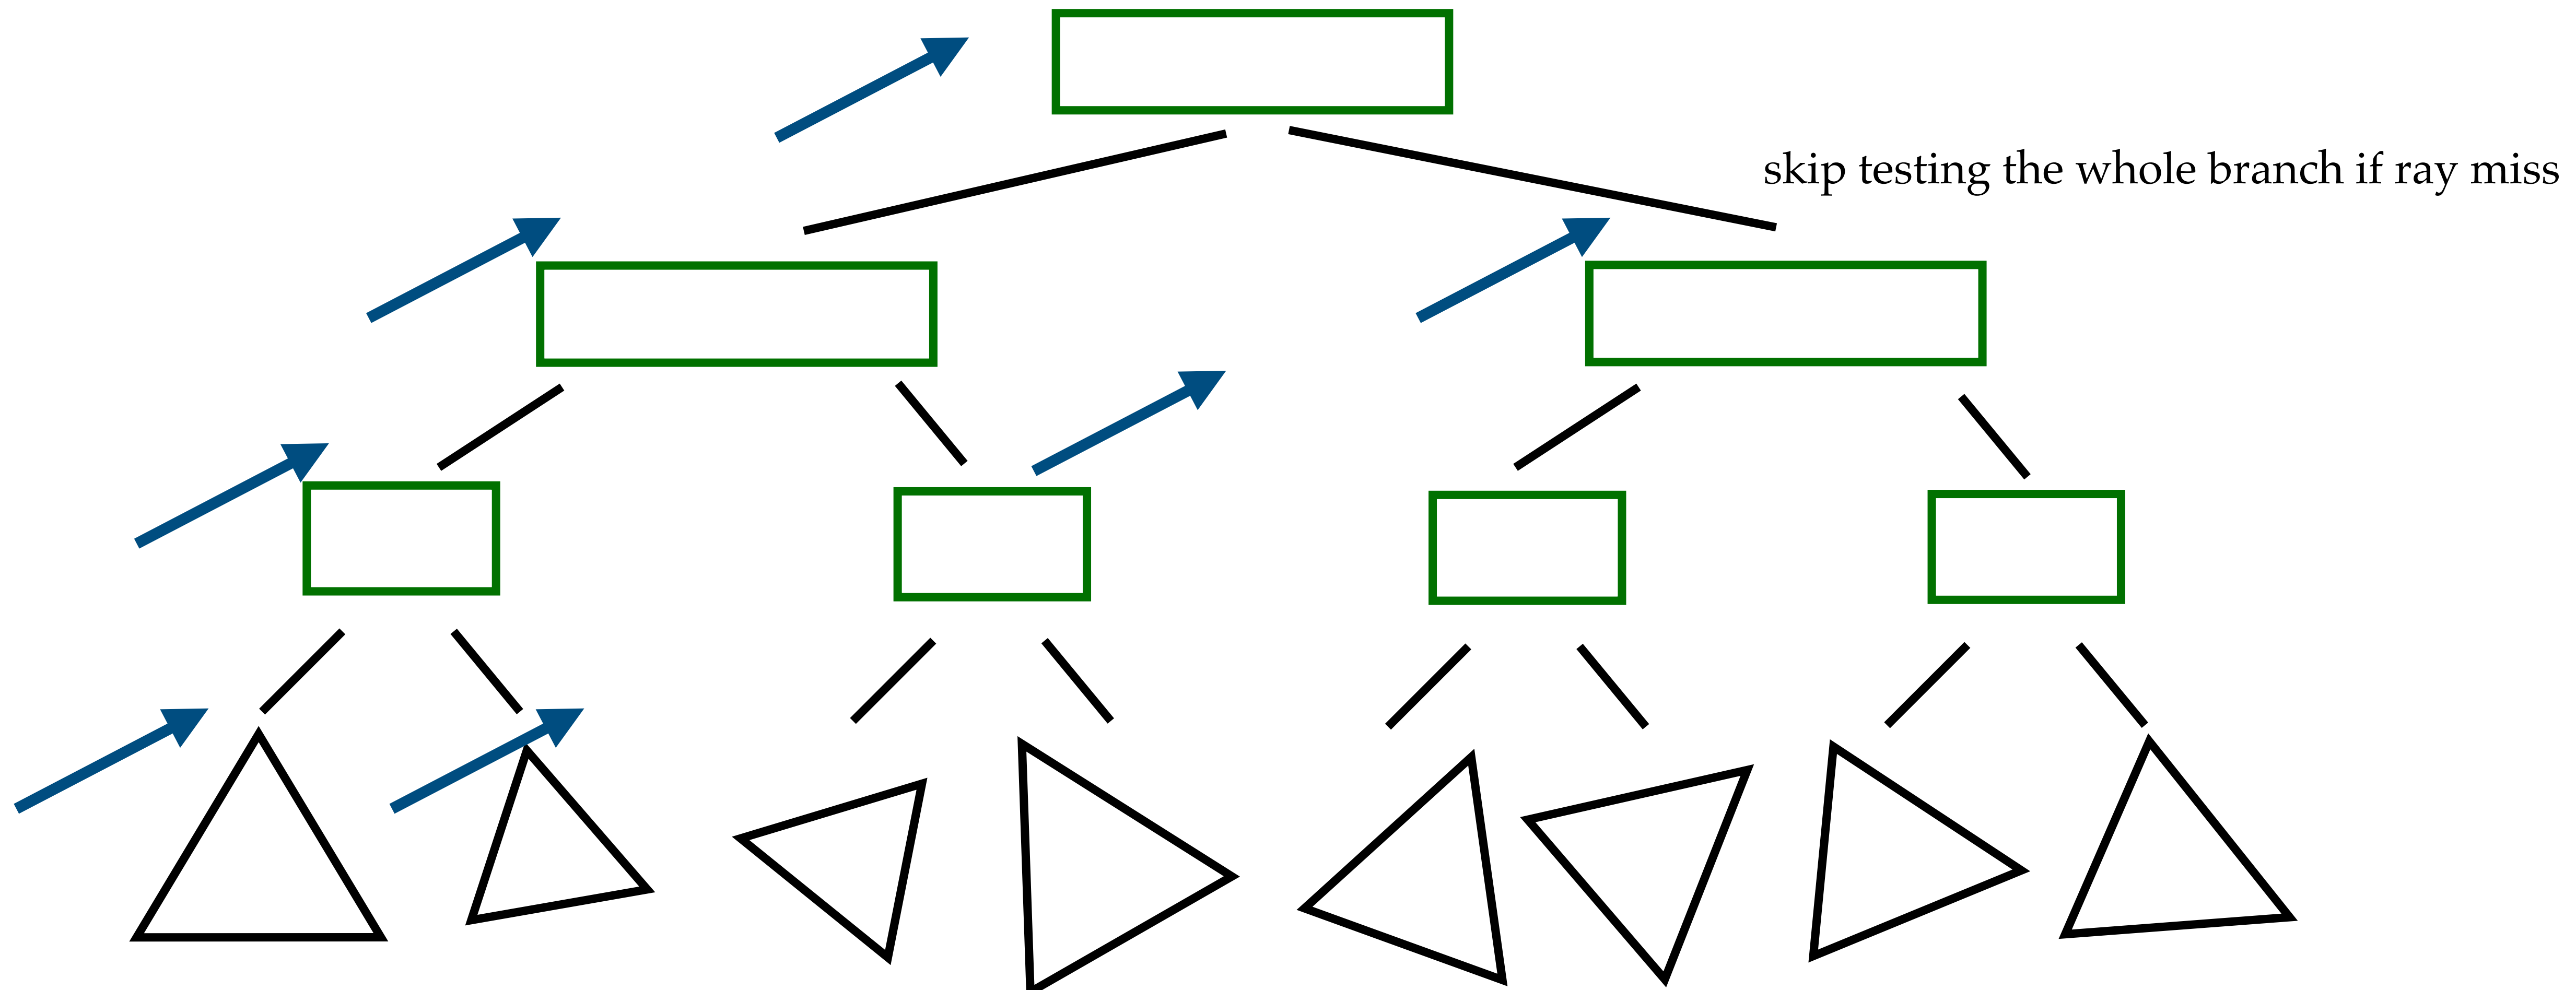
Bounding Volumes Hierarchy

- idea: test a group of triangles at a time by looking at their bounding volumes



Bounding Volumes Hierarchy

- idea: test a group of triangles at a time by looking at their bounding volumes



Decades of research on acceleration structures

- optimized to death, though still an active research area

<h2>Ray Tracing Complex Scenes</h2> <p>Timothy L. Kay James T. Kajiya California Institute of Technology Pasadena, CA 91125</p>	<h2>Space Subdivision for Fast Ray Tracing</h2> <p>Andrew S. Glassner University of North Carolina at Chapel Hill</p>	<h2>Fast Ray Tracing by Ray Classification</h2> <p>James Arvo David Kirk Apollo Computer, Inc. 330 Billerica Road Chelmsford, MA 01824</p>	<h2>Efficiency Issues for Ray Tracing</h2> <p>Brian Smits* University of Utah February 19, 1999</p>
<h2>Filtering, Clustering and Hierarchy Construction: a New Solution for Ray-Tracing Complex Scenes</h2> <p>Frédéric Cazals,¹ George Drettakis,^{1,2} Claude Puech¹</p>	<h2>Dual-Split Trees</h2> <p>Daqi Lin University of Utah</p> <p>Konstantin Shkurko University of Utah</p> <p>Ian Mallett Univers</p>	<h2>Improved Two-Level BVHs using Partial Re-Braiding</h2> <p>Carsten Benthin Intel Corporation</p> <p>Sven Woop Intel Corporation</p> <p>Ingo Wald Intel Corporation</p> <p>Cem Yuksel</p>	<h2>B-KD Trees for Hardware Accelerated Ray Tracing of Dynamic Scenes</h2> <p>Attila T. Áfra Intel Corporation</p> <p>Sven Woop[†], Gerd Marmitt[‡], and Philipp Slusallek[§] Computer Graphics Lab. Saarland University, Germany</p>
<h2>HLBVH: Hierarchical LBBVH Construction for Real-Time Ray Tracing of Dynamic Geometry</h2> <p>J. Pantaleoni¹ and D. Luebke¹ ¹NVIDIA Research</p>	<h2>Two-Level Grids for Ray Tracing on GPUs</h2> <p>Javor Kalojanov¹ and Markus Billeter² and Philipp Slusallek^{1,3}</p>	<h2>Compressed-Leaf Bounding Volume Hierarchies (originally submitted, un-shortened version)</h2> <p>Carsten Benthin Ingo Wald Sven Woop Attila T. Áfra Intel Corporation Intel Corporation</p>	<h2>Wide BVH Traversal with a Short Stack</h2> <p>K. Vaidyanathan S. Woop C. Benthin Intel Corporation</p>
<h2>An N-ary BVH Child Node Sorting Technique for Occlusion Tests</h2> <p>Shinji Ogaki OLM Digital, Inc.</p> <p>Alexandre Derouet-Jourdan OLM Digital, Inc. / JST CREST</p>	<h2>Interactive Rendering with Coherent Ray Tracing</h2> <p>Ingo Wald, Philipp Slusallek, Carsten Benthin, and Markus Wagner Computer Graphics Group, Saarland University</p>	<h2>Fast Parallel Construction of High-Quality Bounding Volume Hierarchies</h2> <p>Tero Karras Timo Aila NVIDIA</p>	<h2>Understanding the Efficiency of Ray Traversal on GPUs</h2> <p>Timo Aila* NVIDIA Research</p> <p>Samuli Laine* NVIDIA Research</p>
<h2>On Quality Metrics of Bounding Volume Hierarchies</h2> <p>Timo Aila Tero Karras Samuli Laine NVIDIA</p>	<h2>MergeTree: A Fast Hardware HLBVH Constructor for Animated Ray Tracing</h2> <p>TIMO VIITANEN, MATIAS KOSKELA, PEKKA JÄÄSKELÄINEN, HEIKKI KULTALA, and JARMO TAKALA, Tampere University of Technology, Finland</p>	<h2>Ray Tracing Lossy Compressed Grid Primitives</h2> <p>Carsten Benthin Karthik Vaidyanathan Sven Woop Intel Corporation</p>	<h2>Using Hardware Ray Transformations for Ray/Primitive Intersections</h2> <p>I. WALD⁽¹⁾ N. MORRICAL^(1,3) S. ZELINSKY⁽²⁾ HUANG⁽²⁾ V. PASCUCCI⁽³⁾, ⁽¹⁾NVIDIA ⁽²⁾</p>

Lajolla uses the Embree library

- highly-optimized ray intersection routines
 - used by almost all production ray tracers

```
/* Intersects a single ray with the scene. */  
void rtcIntersect1(RTCScene scene, RTCIntersectContext* context, RTCRayHit* rayhit);
```

Embree: A Kernel Framework for Efficient CPU Ray Tracing

Ingo Wald † Sven Woop † Carsten Benthin † Gregory S. Johnson † Manfred Ernst ‡
Intel Corporation




Figure 1: Images produced by renderers which use the open source Embree ray tracing kernels. These scenes are computationally challenging due to complex geometry and spatially incoherent secondary rays. From left to right: The White Room model by Jay Hardy rendered in Autodesk RapidRT, a car model rendered in the Embree path tracer, a scene from the DreamWorks Animation movie “Peabody & Sherman” rendered with a prototype path tracer, and the Imperial Crown of Austria model by Martin Lubich rendered in the Embree path tracer.

Intersection function in lajolla


```
/// Intersect a ray with a scene. If the ray doesn't hit anything,  
/// returns an invalid optional output.  
std::optional<PathVertex> intersect(const Scene &scene,  
                                   const Ray &ray,  
                                   const RayDifferential &ray_diff = RayDifferential{});
```


Intersection function in lajolla

```
/// Intersect a ray with a scene. If the ray doesn't hit anything,  
/// returns an invalid optional output.  
std::optional<PathVertex> intersect(const Scene &scene,  
                                   const Ray &ray,  
                                   const RayDifferential &ray_diff = RayDifferential{});  
  
struct PathVertex {  
    Vector3 position;  
    Vector3 geometry_normal; // always face at the same direction at shading_frame.n  
    Frame shading_frame;  
    Vector2 st; // A 2D parametrization of the surface. Irrelavant to UV mapping.  
               // for triangle this is the barycentric coordinates, which we use  
               // for interpolating the uv map.  
    Vector2 uv; // The actual UV we use for texture fetching.  
             // For texture filtering, stores approximatedly min(abs(du/dx), abs(dv/dx), abs(du/dy), abs(dv/dy))  
    Real uv_screen_size;  
    Real mean_curvature; // For ray differential propagation.  
    Real ray_radius; // For ray differential propagation.  
    int shape_id = -1;  
    int primitive_id = -1; // For triangle meshes. This indicates which triangle it hits.  
    int material_id = -1;  
  
    // ...  
};
```


Smallpt: constant color across the surface

```
1. Sphere spheres[] = {//Scene: radius, position, emission, color, material
2.   Sphere(1e5, Vec( 1e5+1, 40.8, 81.6), Vec(), Vec(.75, .25, .25), DIFF), //Left
3.   Sphere(1e5, Vec(-1e5+99, 40.8, 81.6), Vec(), Vec(.25, .25, .75), DIFF), //Right
4.   Sphere(1e5, Vec(50, 40.8, 1e5), Vec(), Vec(.75, .75, .75), DIFF), //Back
5.   Sphere(1e5, Vec(50, 40.8, -1e5+170), Vec(), Vec(), DIFF), //Frnt
6.   Sphere(1e5, Vec(50, 1e5, 81.6), Vec(), Vec(.75, .75, .75), DIFF), //Botm
7.   Sphere(1e5, Vec(50, -1e5+81.6, 81.6), Vec(), Vec(.75, .75, .75), DIFF), //Top
8.   Sphere(16.5, Vec(27, 16.5, 47), Vec(), Vec(1, 1, 1)*.999, SPEC), //Mirr
9.   Sphere(16.5, Vec(73, 16.5, 78), Vec(), Vec(1, 1, 1)*.999, REFR), //Glas
10.  Sphere(600, Vec(50, 681.6-.27, 81.6), Vec(12, 12, 12), Vec(), DIFF) //Lite
11. };
```



Real-world surfaces are colorful!




photo from K.C. Alfred

<https://www.sandiegouniontribune.com/sdut-snake-path-2012jan20-htmlstory.html>


An option: assign a color to each triangle

quiz: what are the pros and cons?





An option: assign a color to each triangle

- pros
 - simple
 - easy to edit (just paint on triangles)
- cons
 - couples geometric complexity with color complexity
 - hard to **filter**
 - more on this later



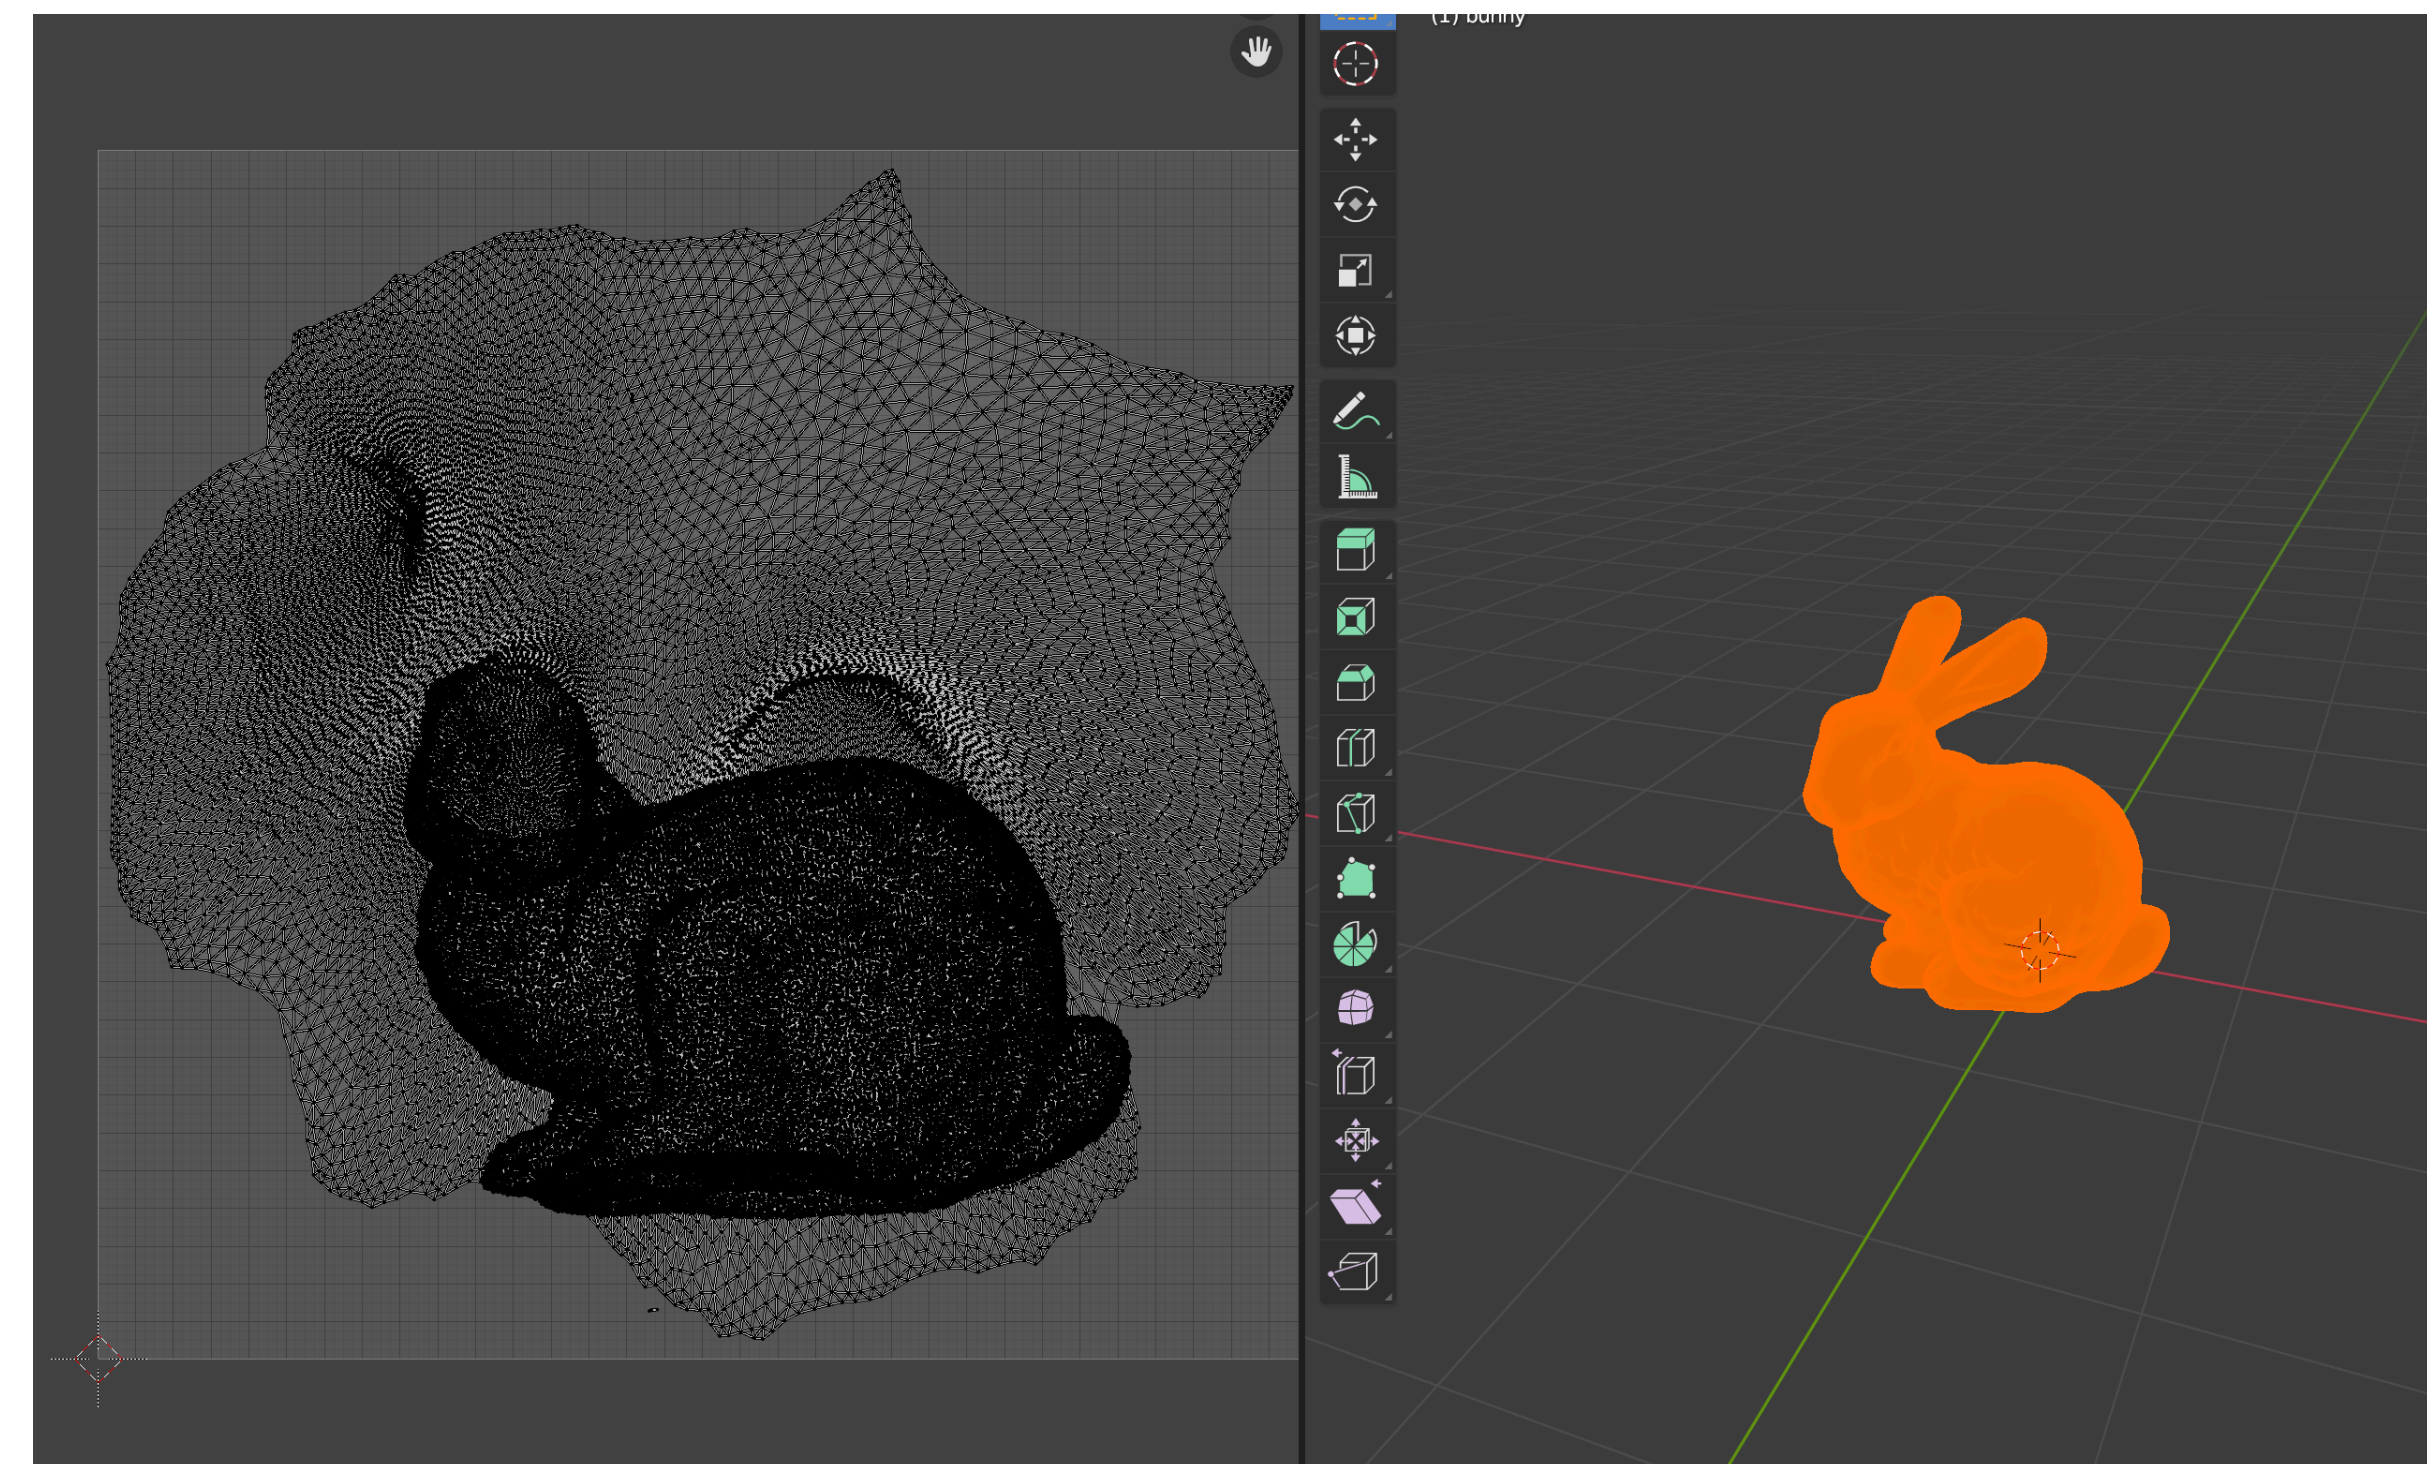
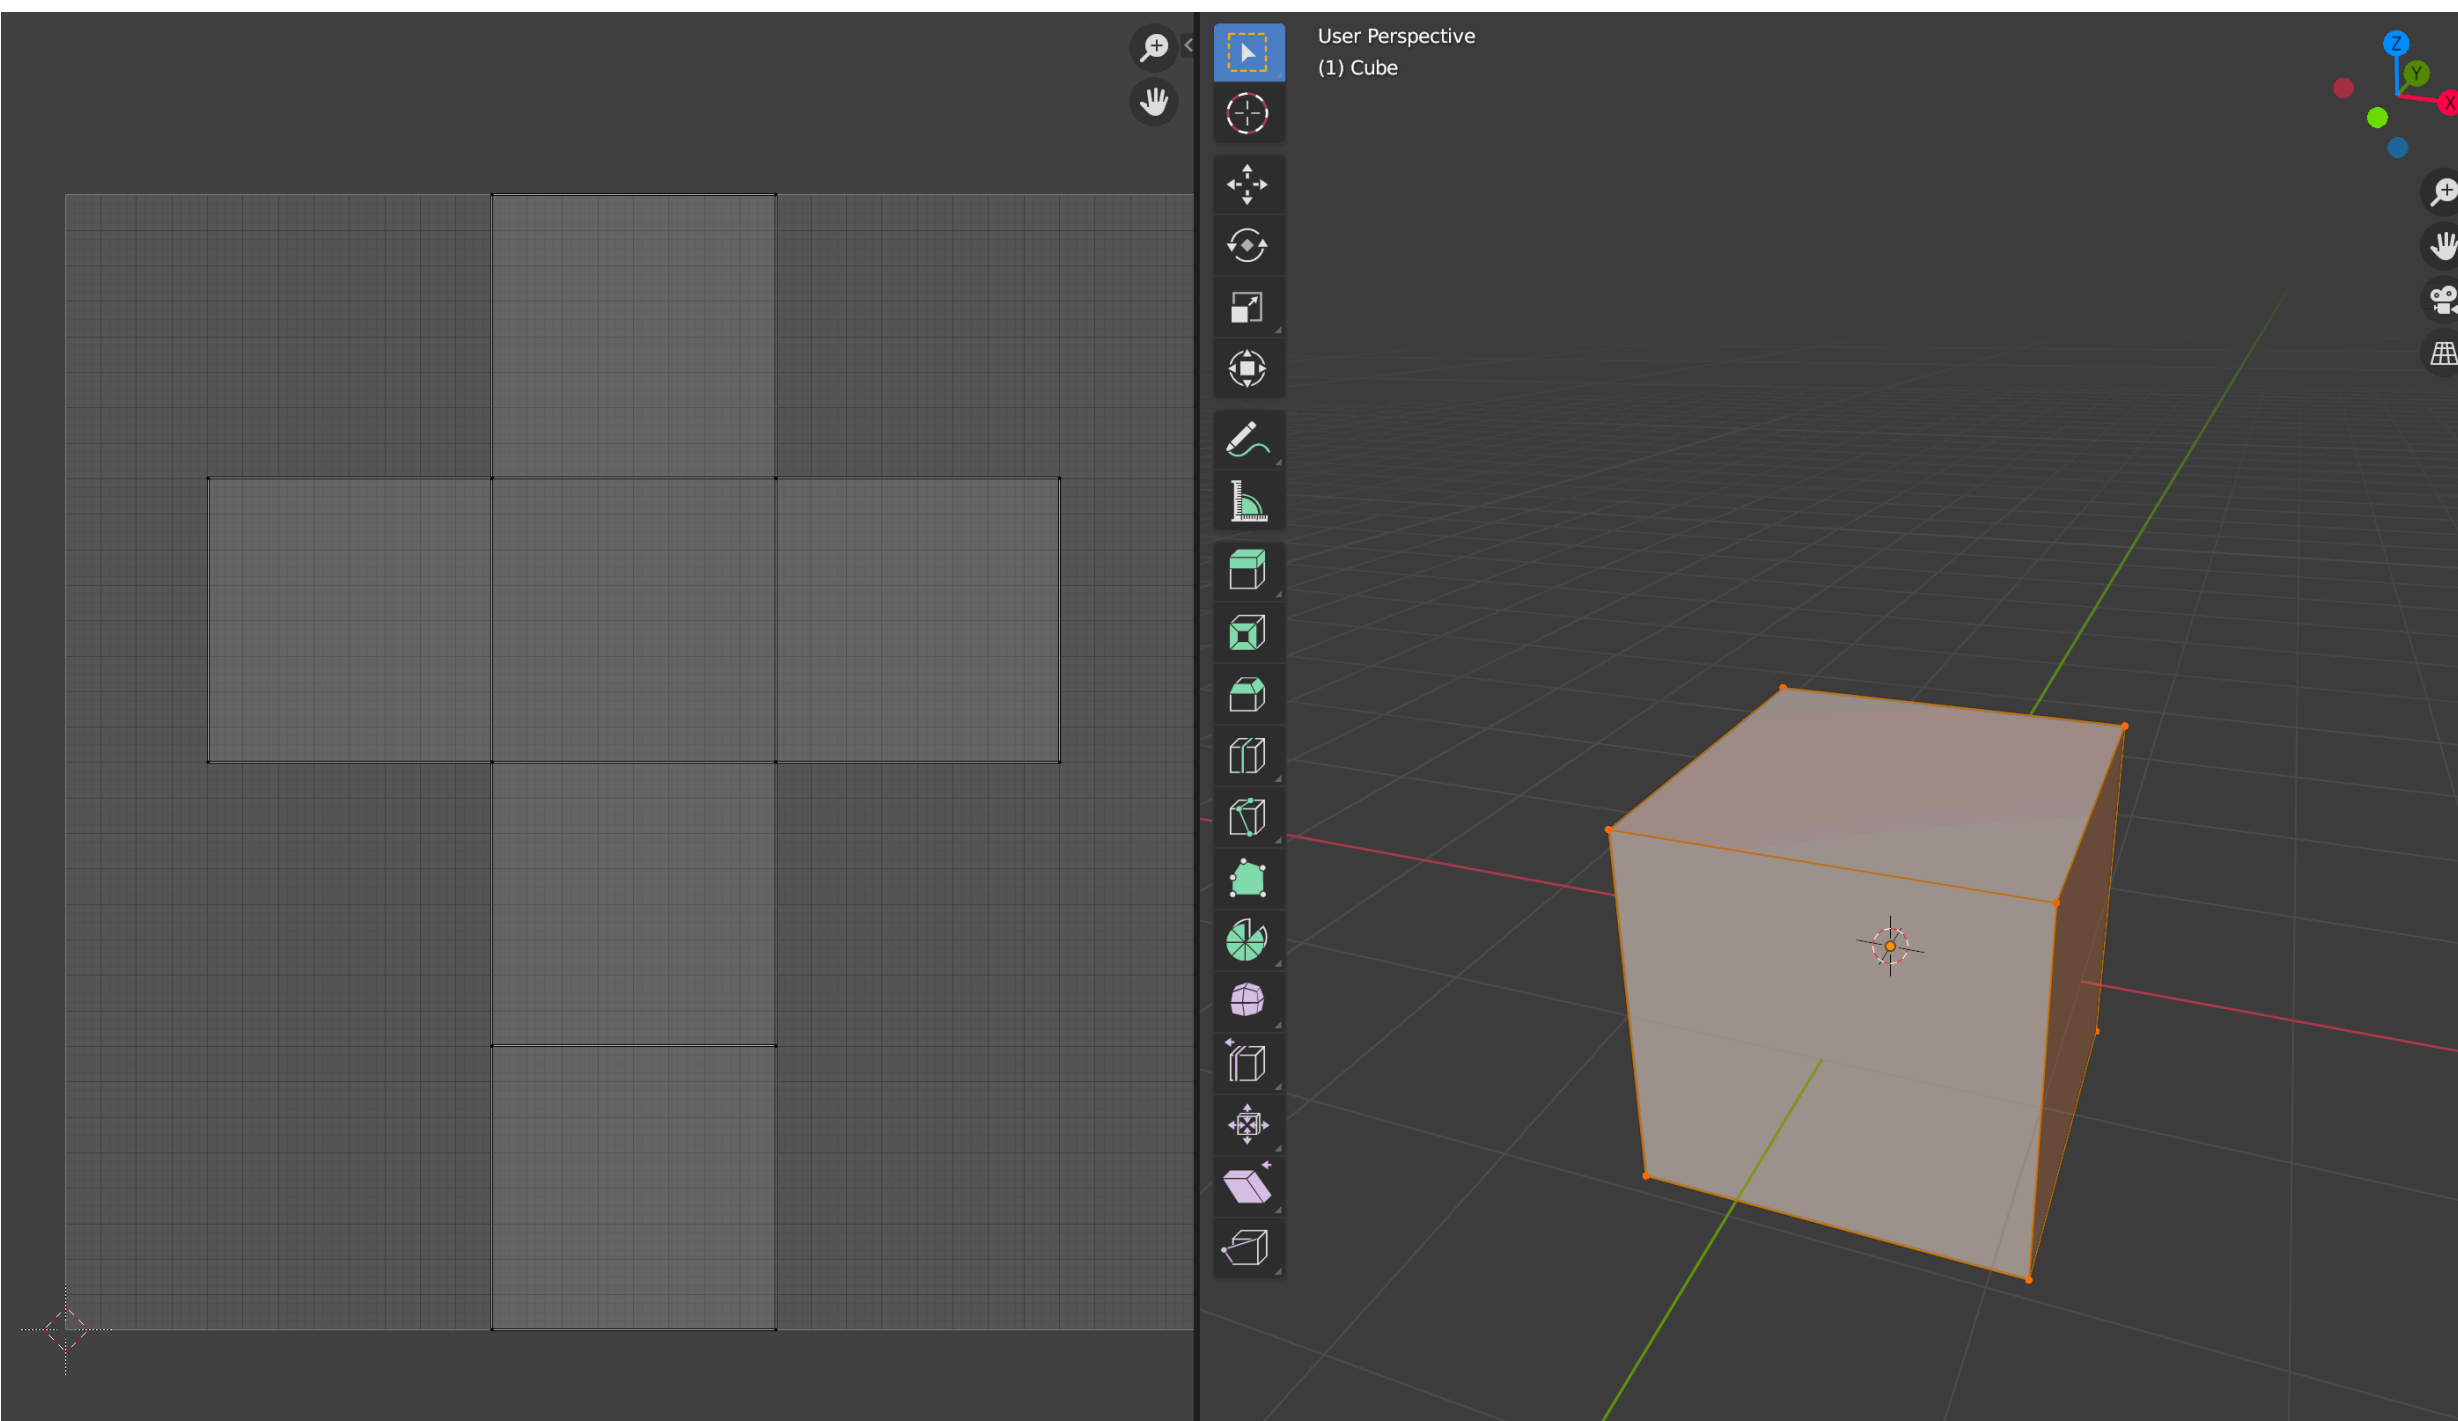
In practice: texture mapping

- assign a “UV” 2D vector to each point on the surface



UV mapping

- “unwrap” a surface and map it to a 2D square
- automatic UV mapping is an active research area



In lajolla


```
struct TriangleMesh : public ShapeBase {  
    std::vector<Vector3> positions;  
    std::vector<Vector3i> indices;  
    std::vector<Vector3> normals;  
    std::vector<Vector2> uvs;  
    // ...  
};
```

uv per-vertex

```
// Barycentric coordinates are stored in vertex.st  
Vector2 uv = (1 - vertex.st[0] - vertex.st[1]) * uvs[0] +  
             vertex.st[0] * uvs[1] +  
             vertex.st[1] * uvs[2];
```

```
struct PathVertex {  
    // ...  
    Vector2 st; // A 2D parametrization of the surface. Irrelavant to UV mapping.  
                // for triangle this is the barycentric coordinates, which we use  
                // for interpolating the uv map.  
    Vector2 uv; // The actual UV we use for texture fetching.  
    // ...  
};
```


Obtain UV by interpolating values from vertices



$$uv = (1 - b_1 - b_2)uv_0 + b_1uv_1 + b_2uv_2$$

Texture mapping

quiz: what are the pros and cons?



Texture mapping: pros and cons

- pros
 - different sampling rates for geometry and color
 - much easier to filter
- cons
 - uv mapping is hard





Texture mapping: pros and cons


- pros
 - different sampling rates for geometry and color
 - much easier to filter
- cons
 - uv mapping is hard




A pixel can cover a large region in the texture



Fail to account for all texels in the region
can lead to noise / aliasing

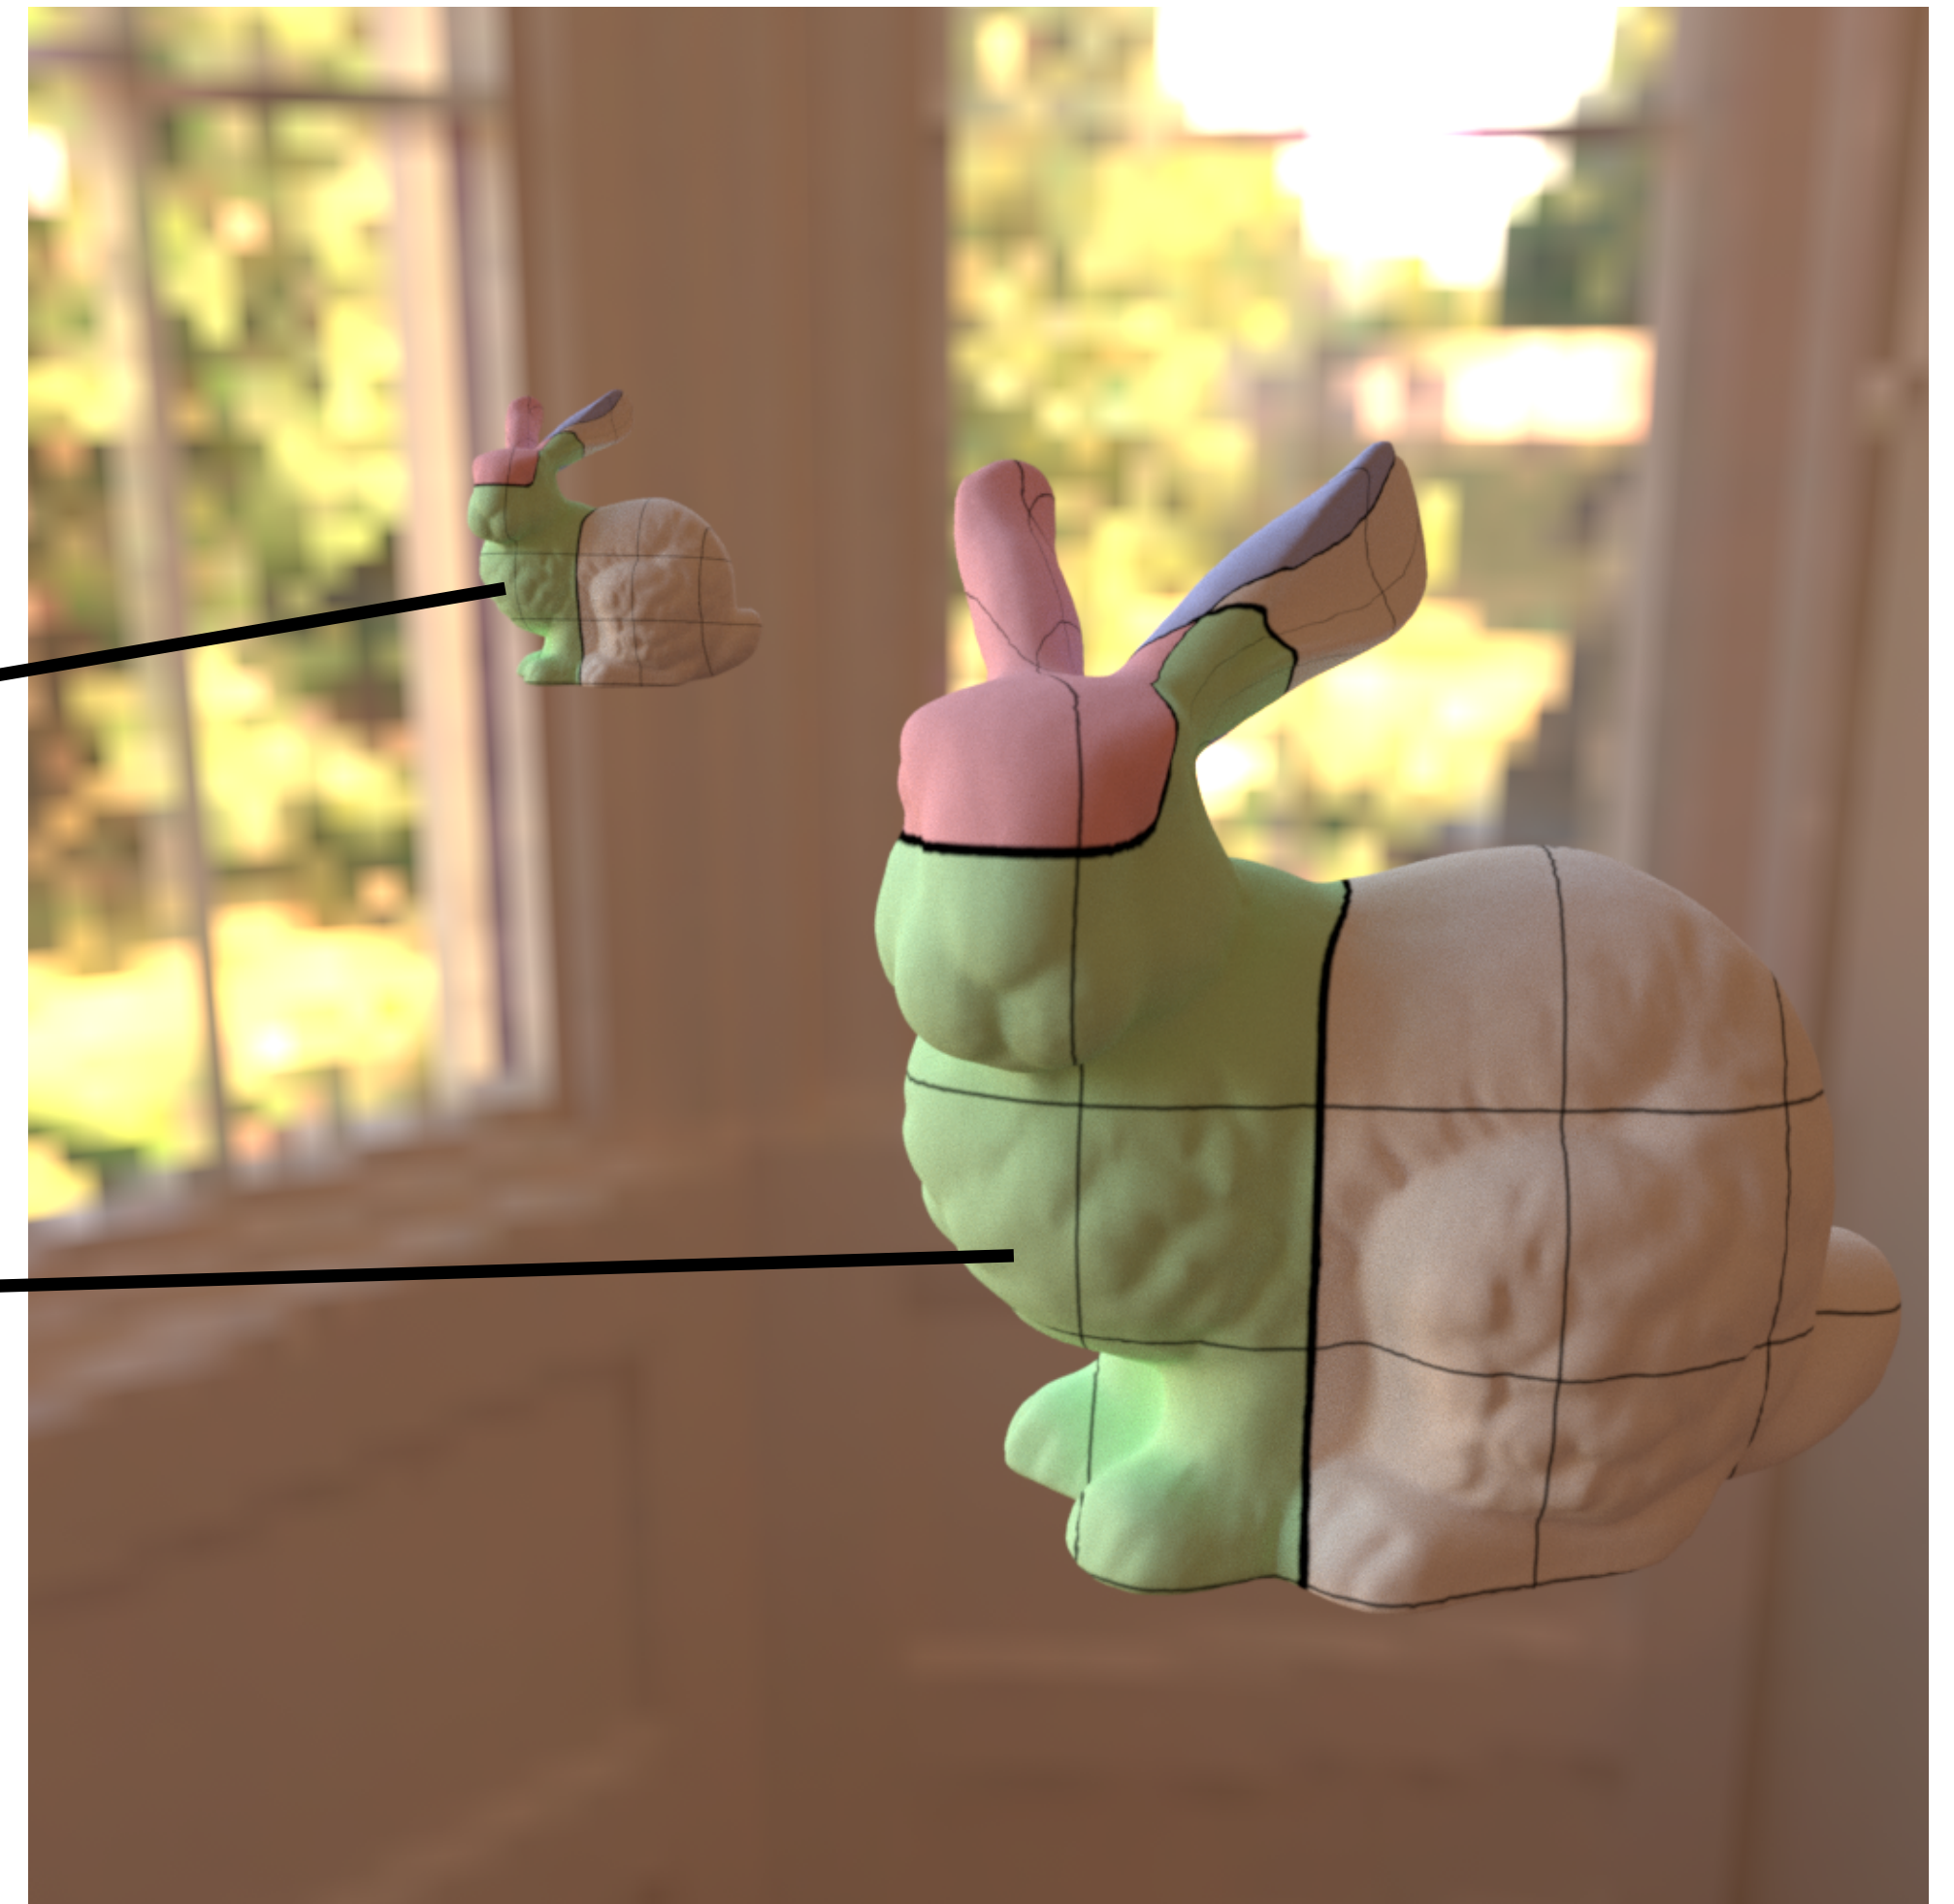



no filtering



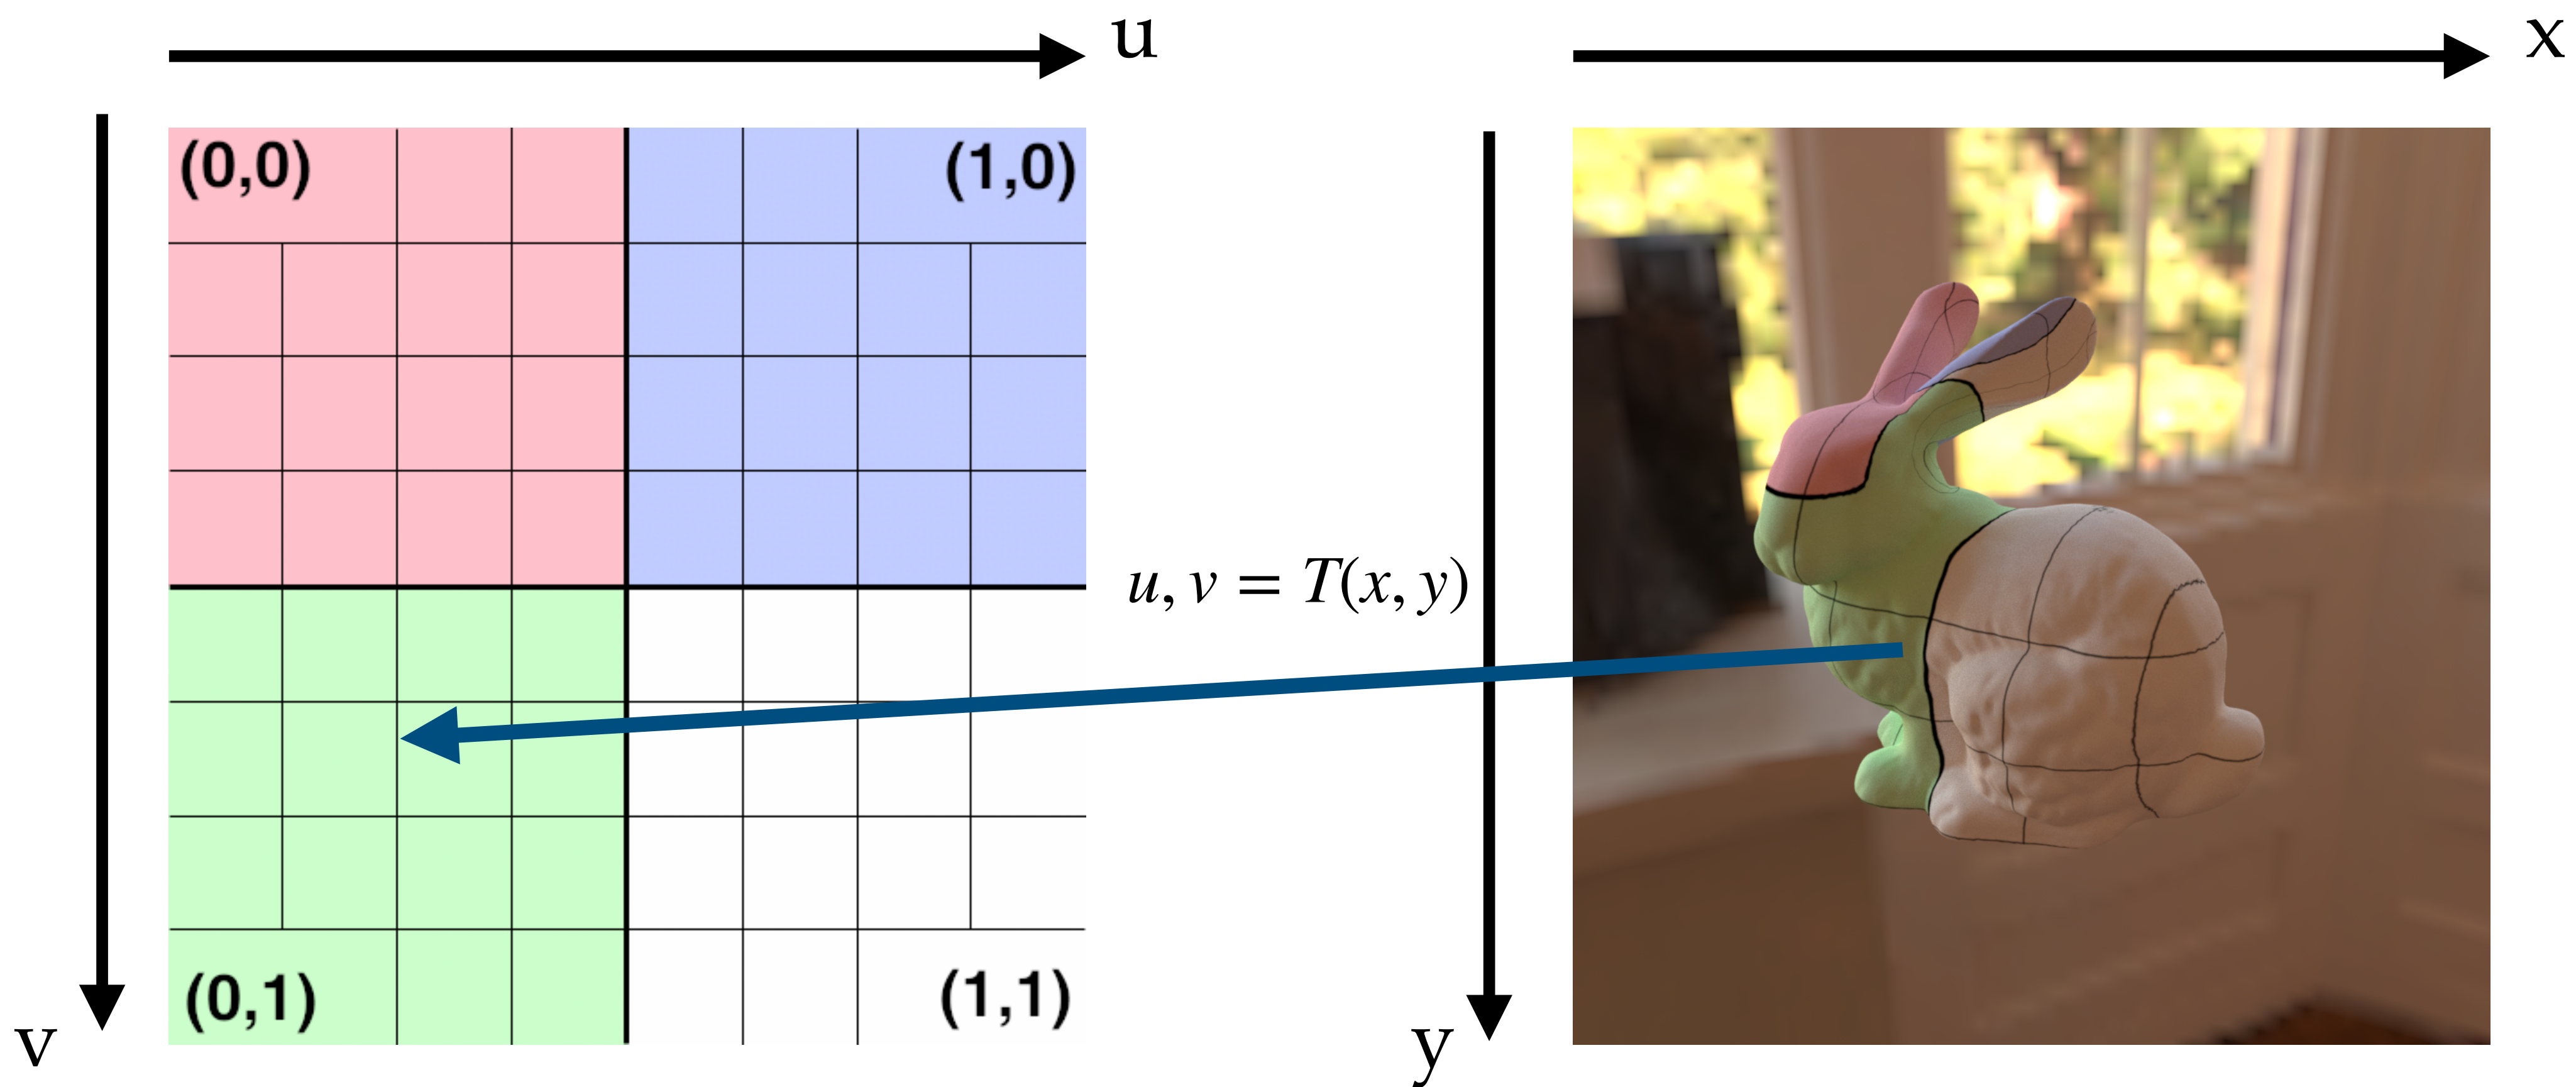
with filtering

Goal: average all texture values inside the region




The filtering region is determined by the mapping between image space and texture space

problem: T can be complicated



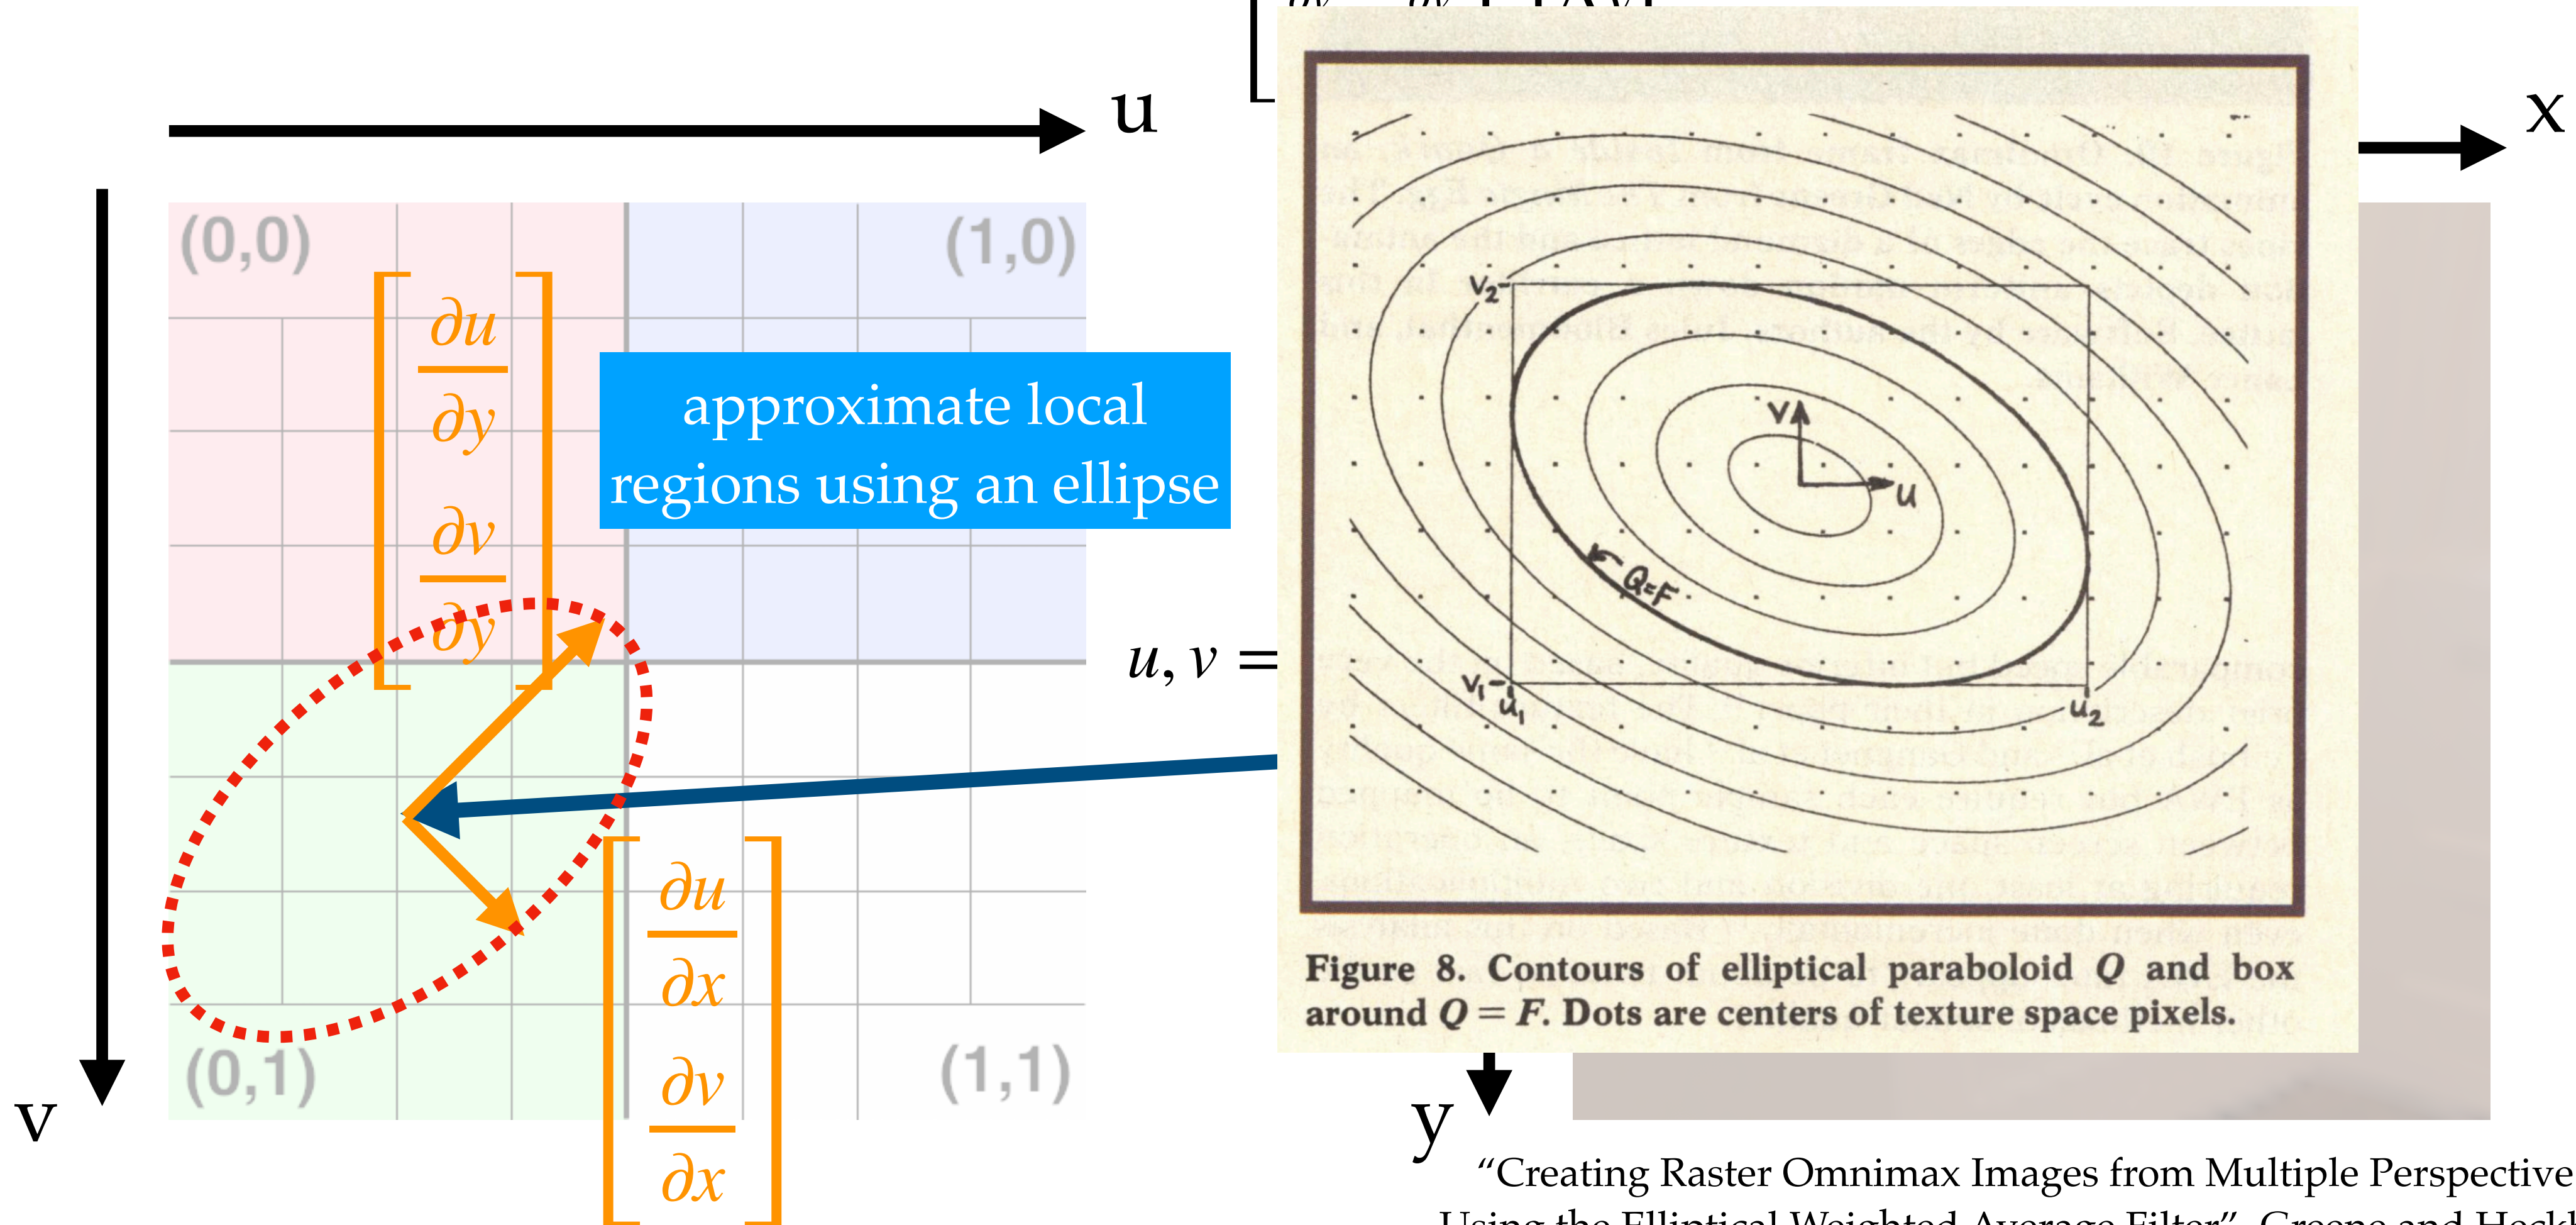
We can approximate the local mapping using first-order Taylor expansion

$$u, v \approx T(x_0, y_0) + \begin{bmatrix} \frac{\partial u}{\partial x} & \frac{\partial u}{\partial y} \\ \frac{\partial v}{\partial x} & \frac{\partial v}{\partial y} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}$$

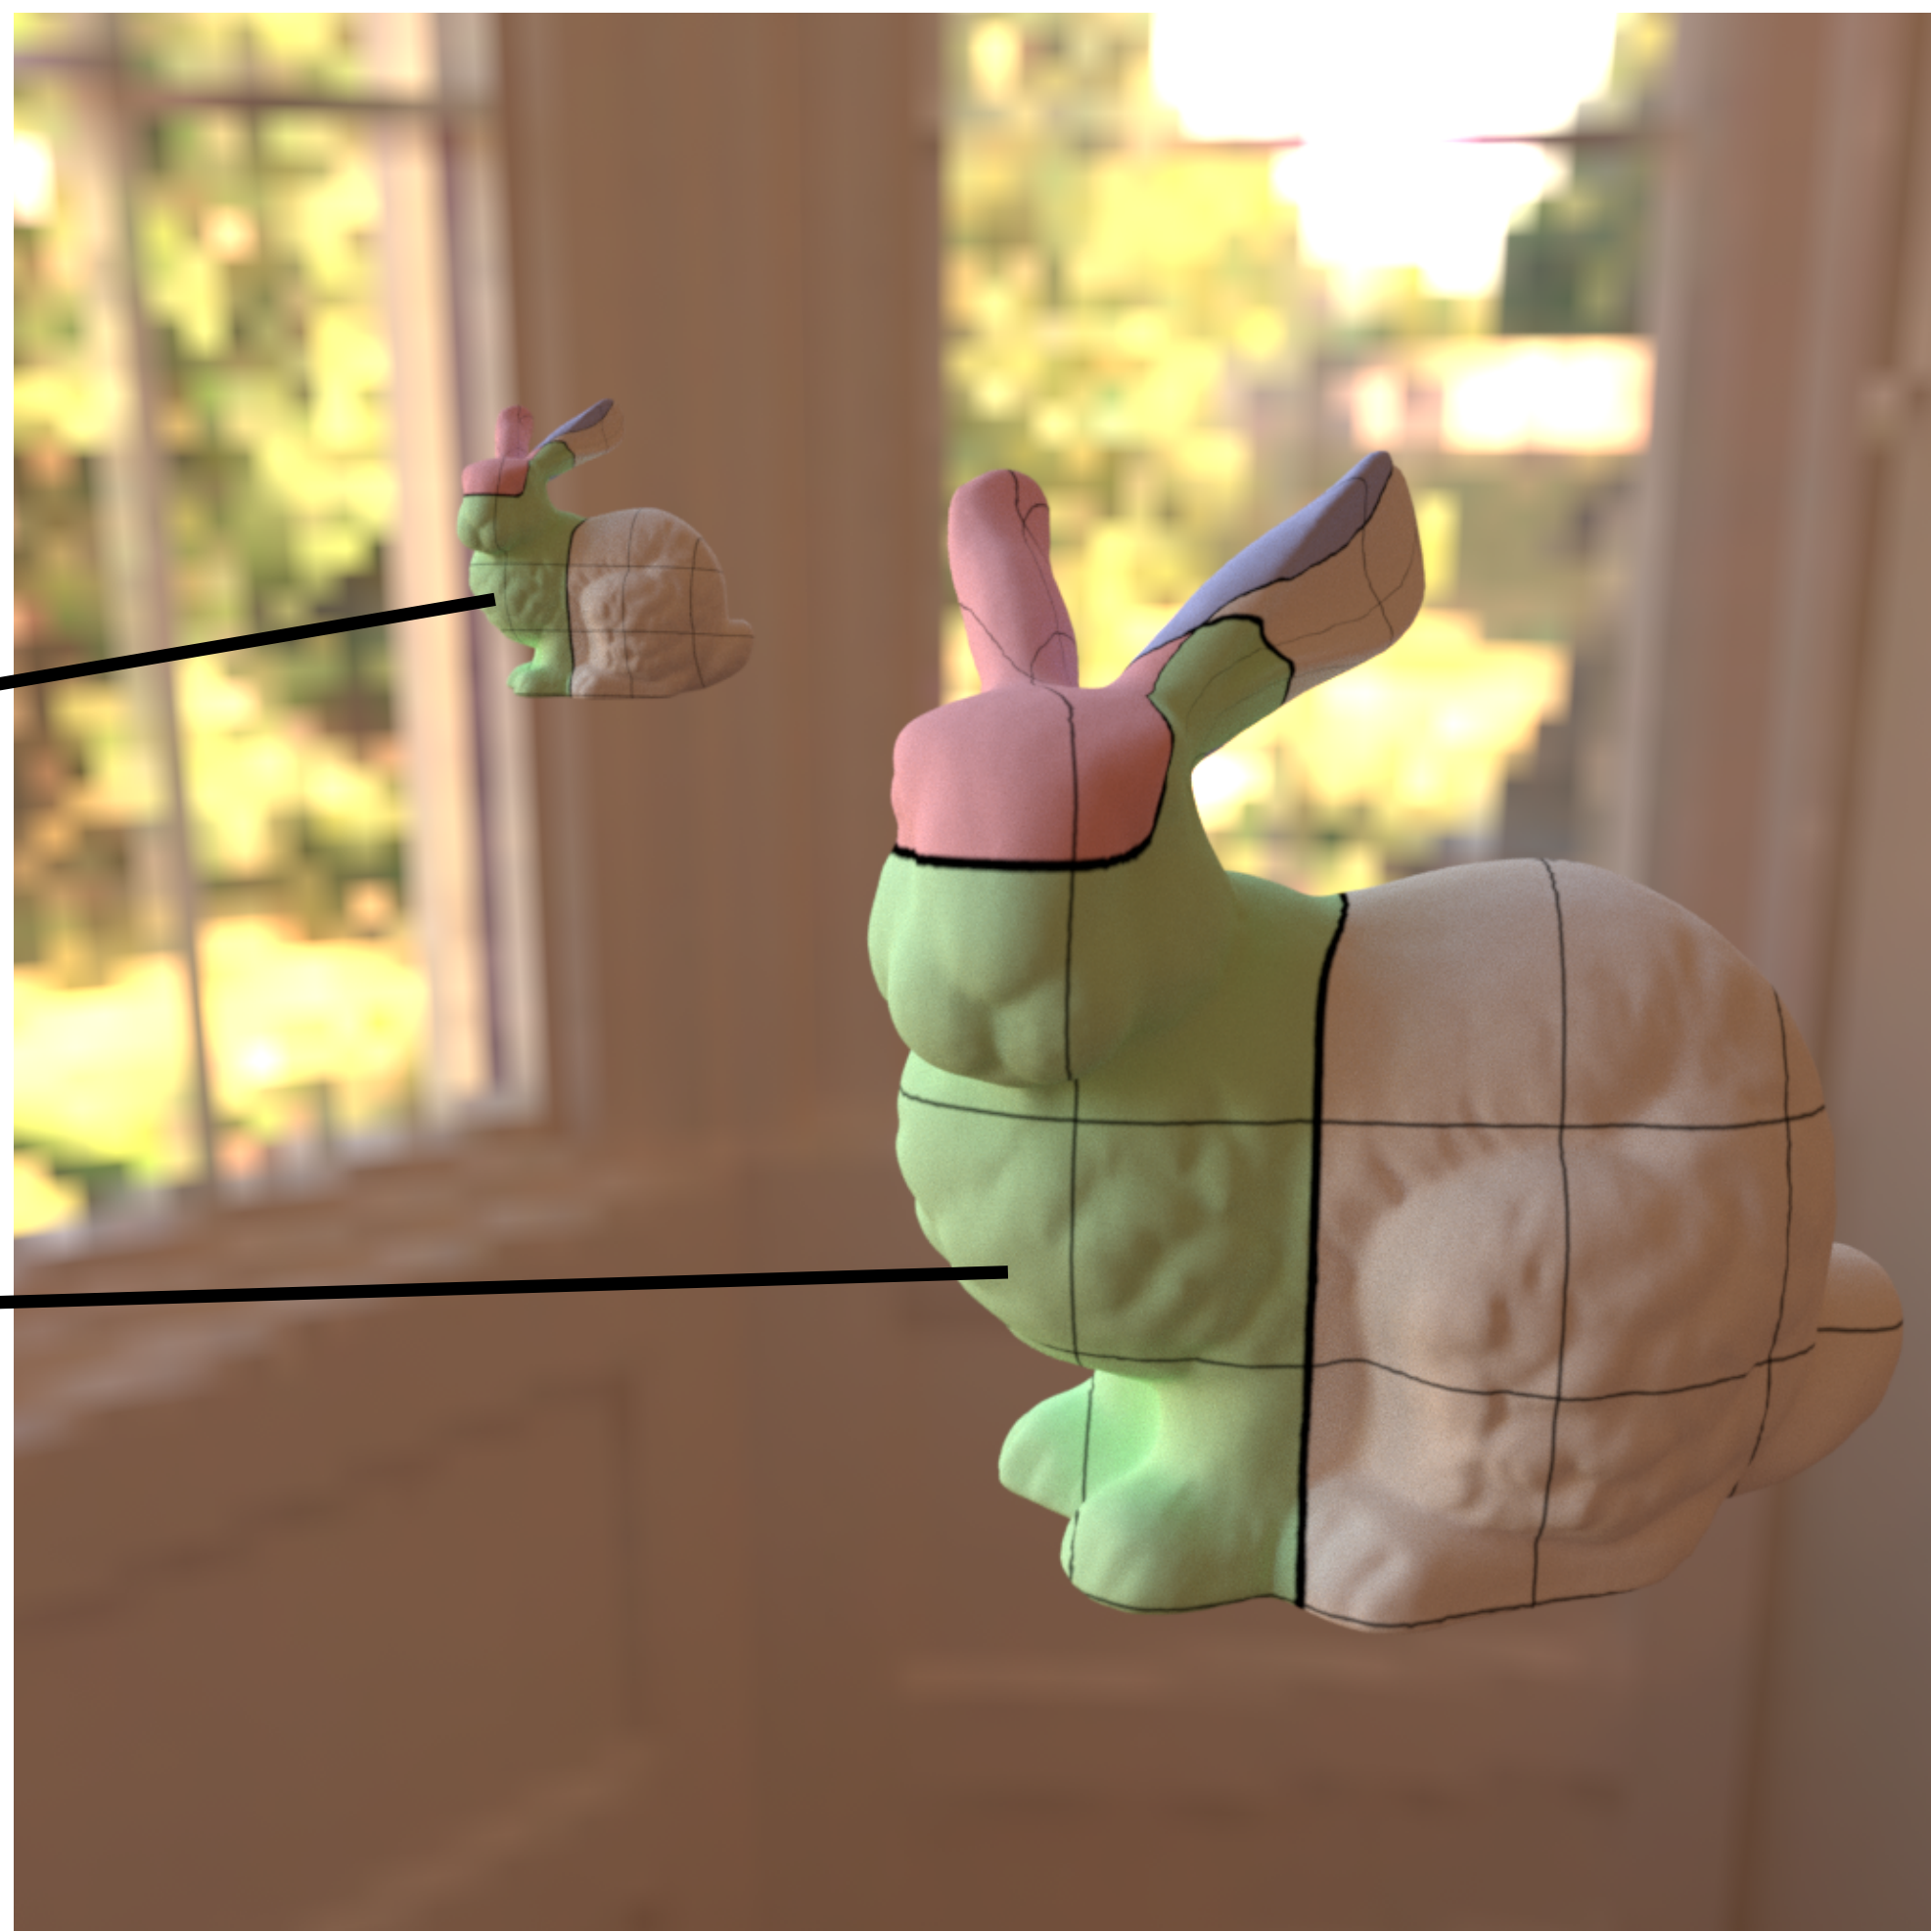



We can approximate the local mapping using first-order Taylor expansion

$$u, v \approx T(x_0, y_0) + \begin{bmatrix} \frac{\partial u}{\partial x} & \frac{\partial u}{\partial y} \\ \frac{\partial v}{\partial x} & \frac{\partial v}{\partial y} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}$$

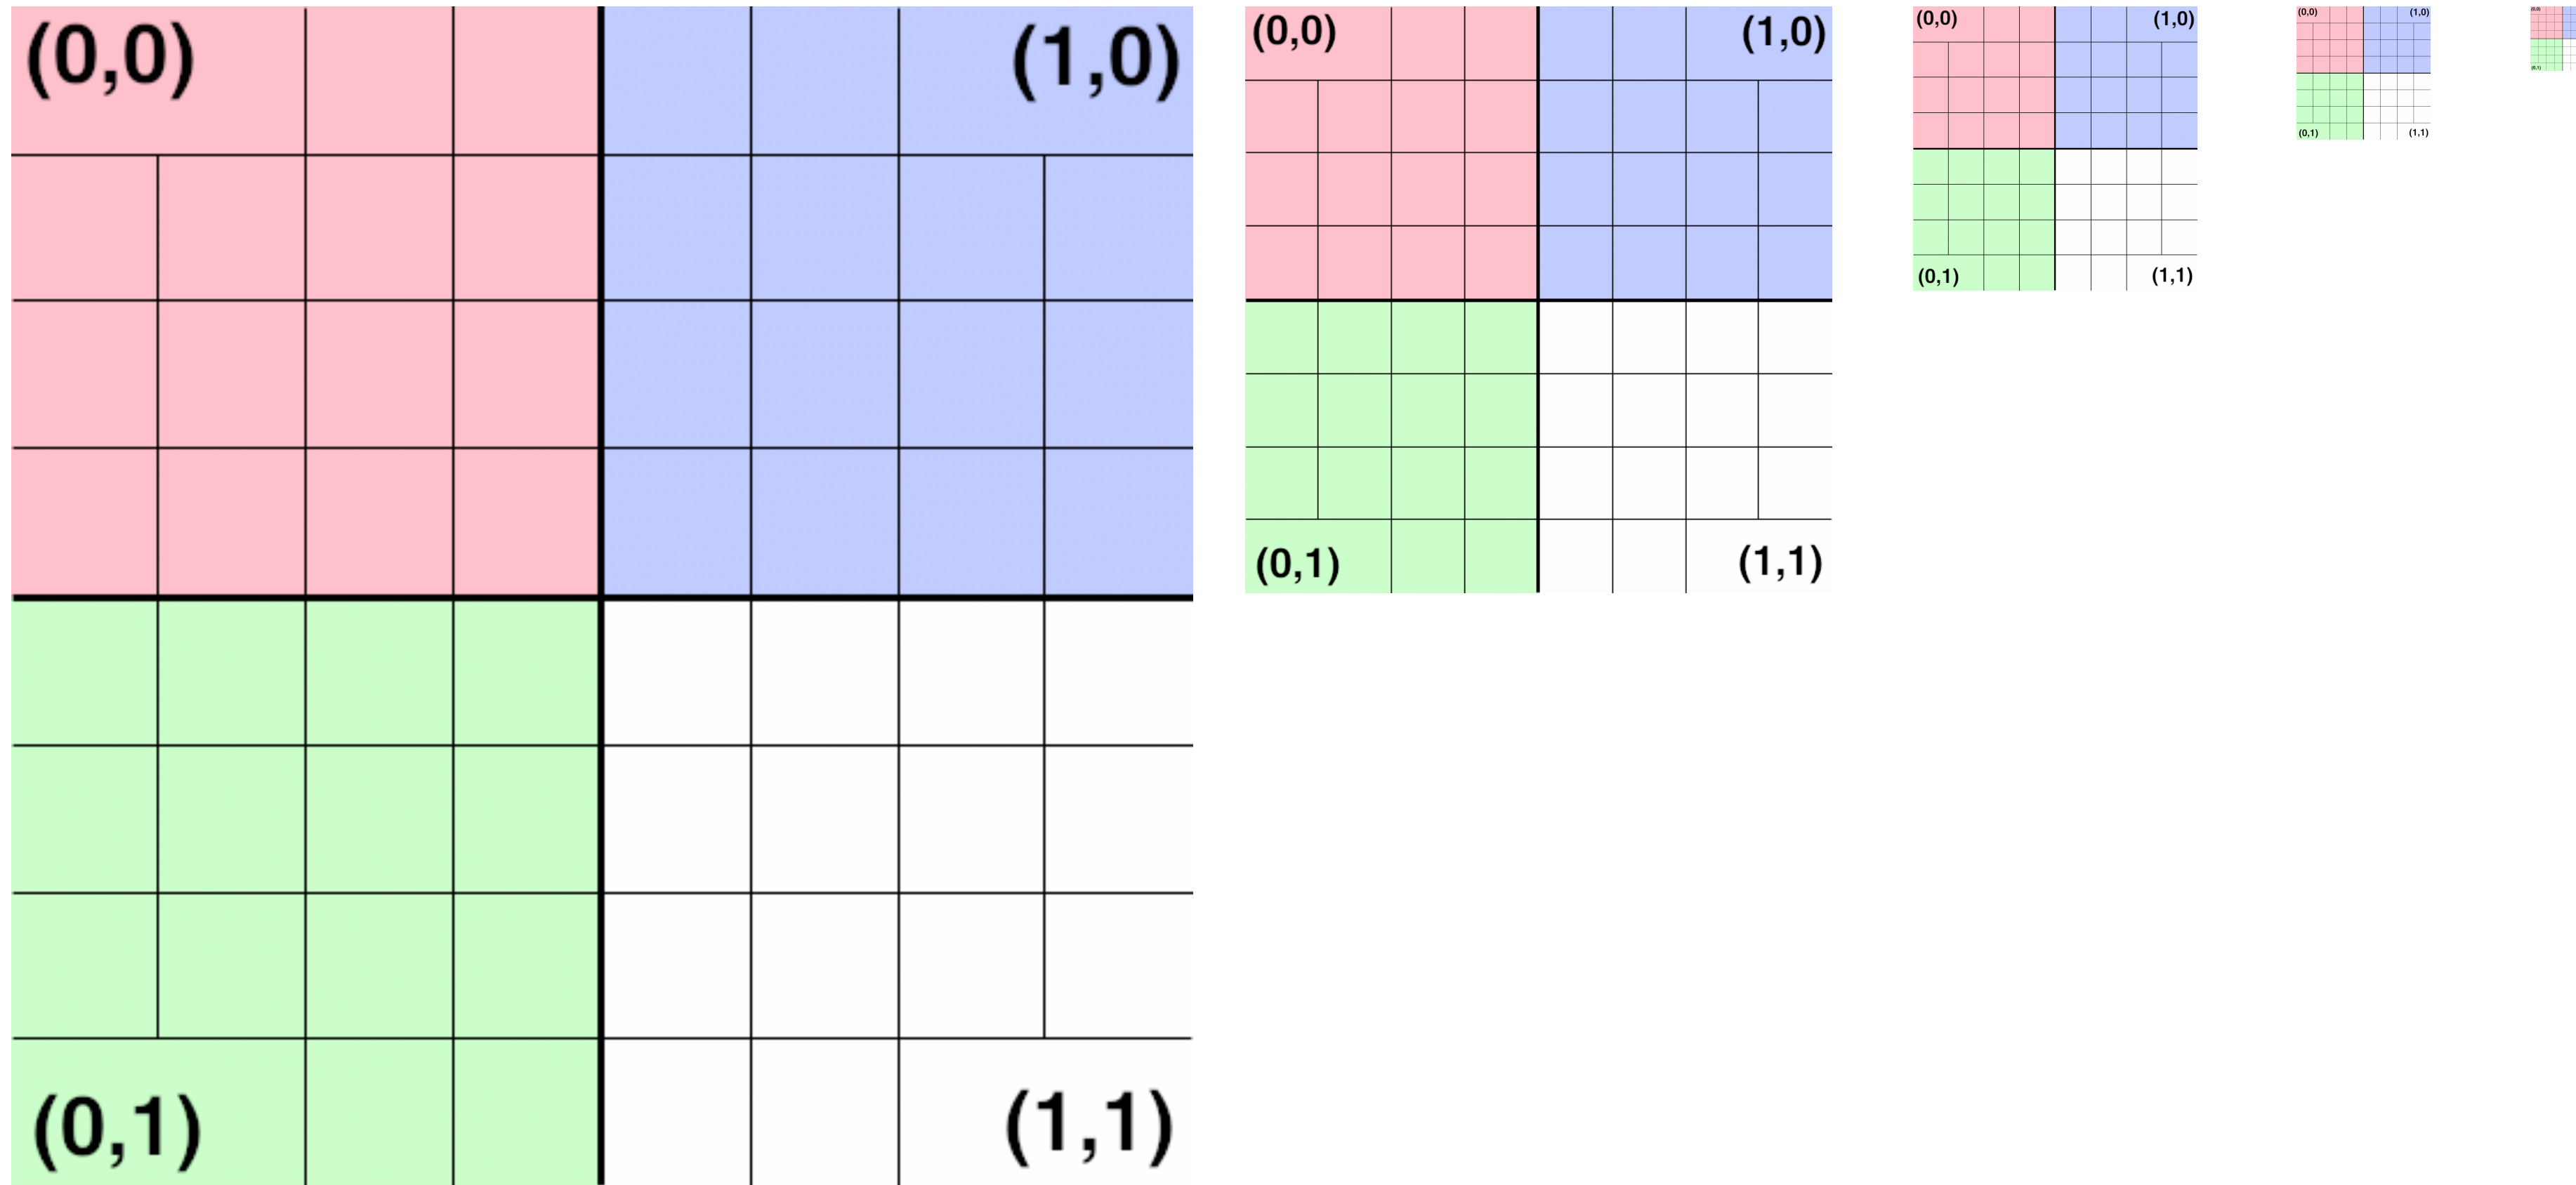


Goal: average all texture values inside the region




Downsample the texture for fast average

- usually called “mipmapping”




Elliptical weighted averaging algorithm



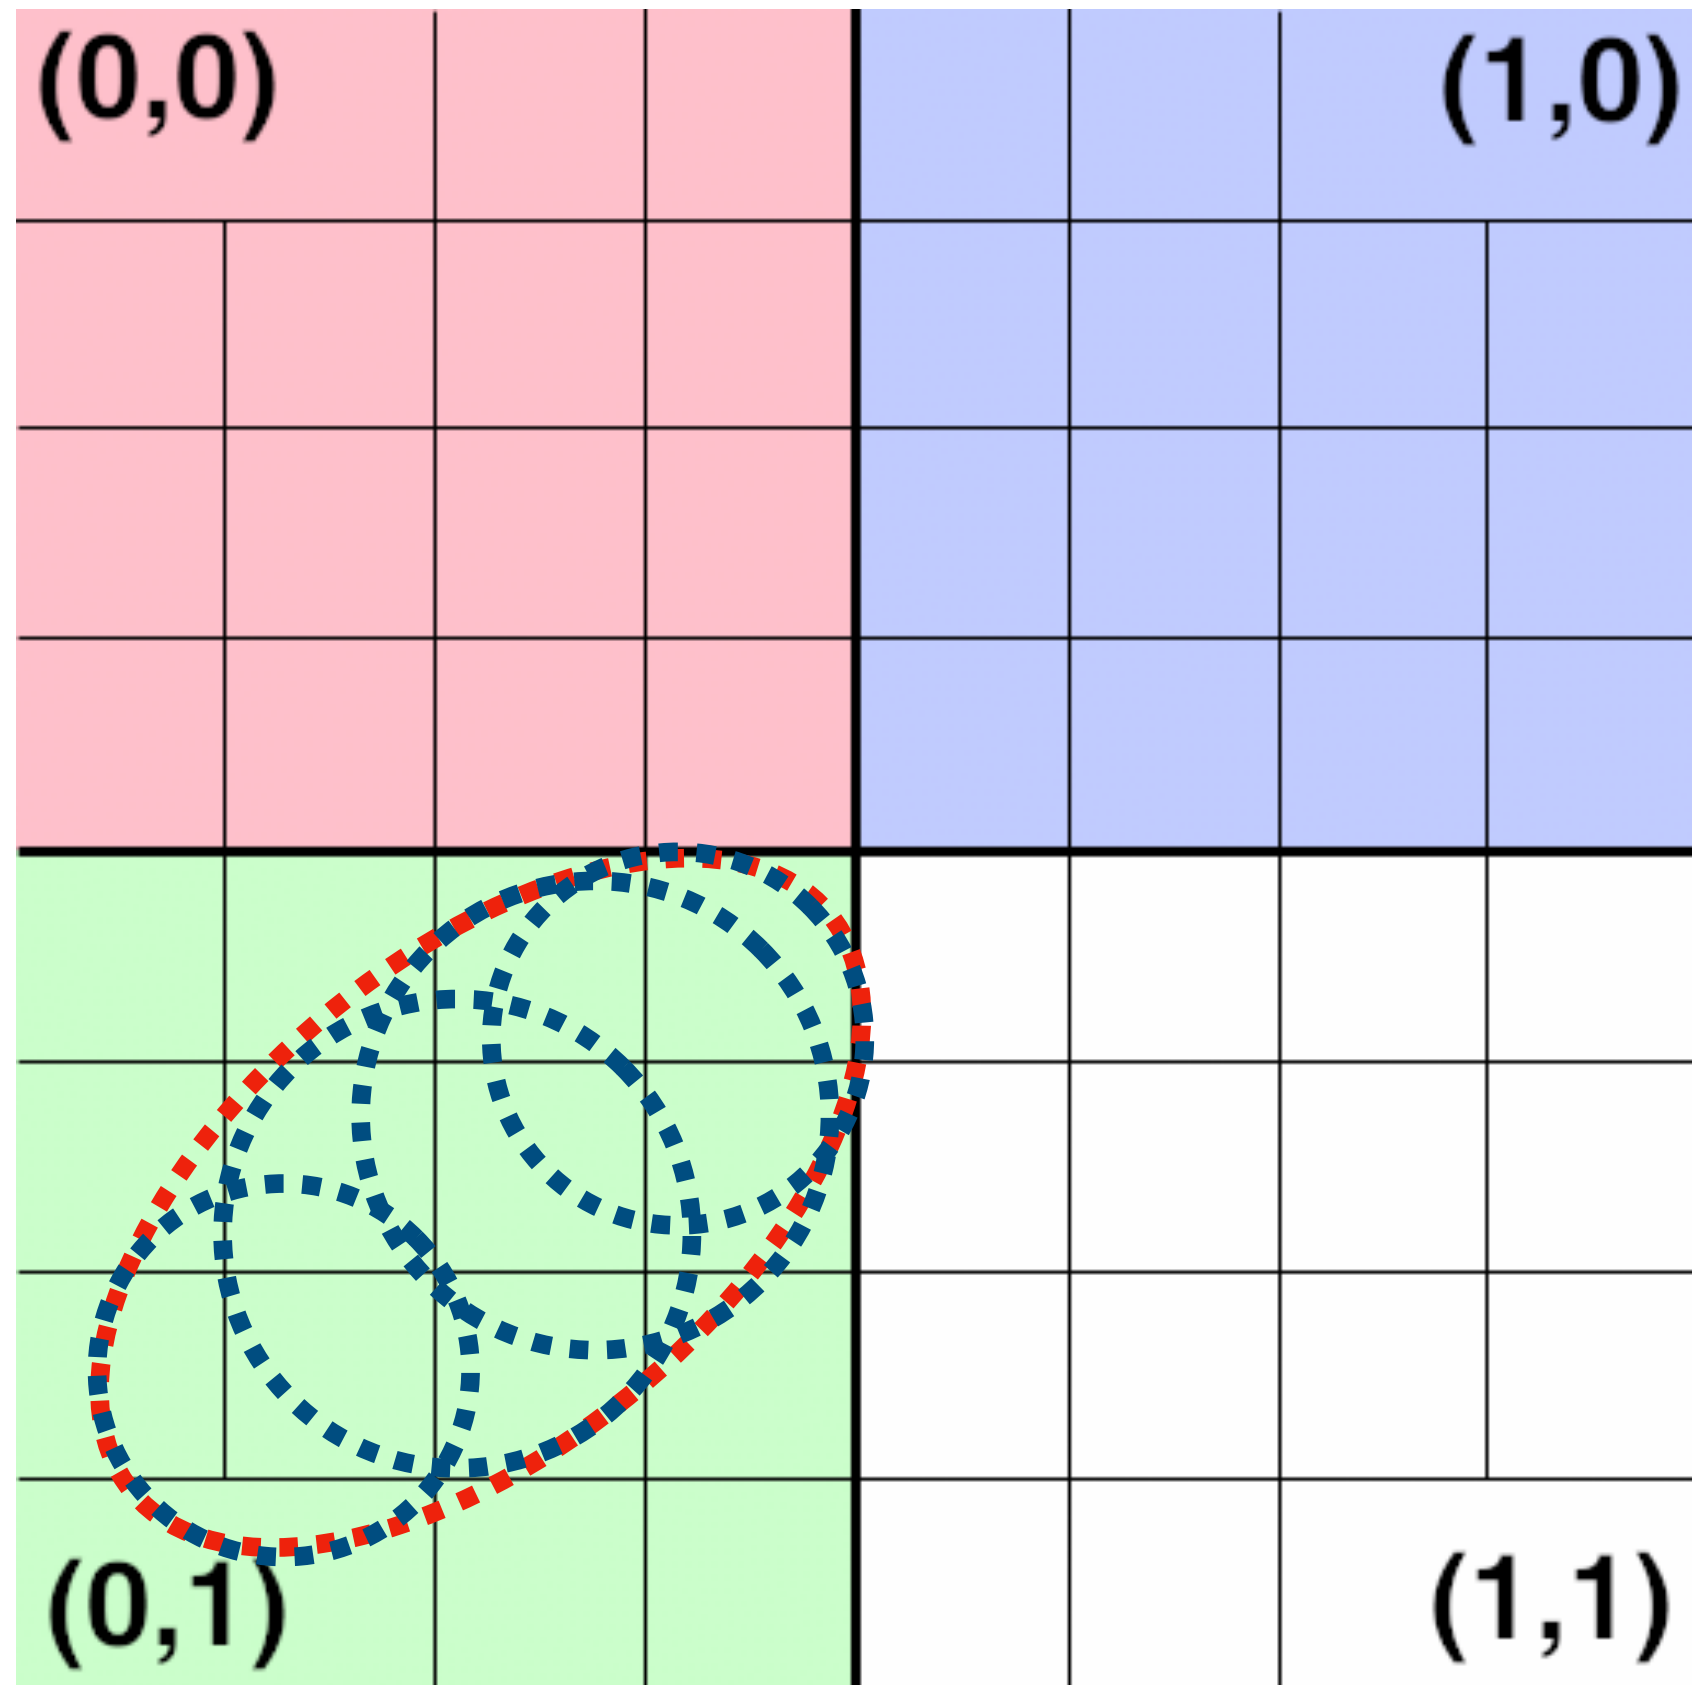
Elliptical weighted averaging algorithm

1. approximate the ellipse with circles




Elliptical weighted averaging algorithm

1. approximate the ellipse with circles




2. find two appropriate mipmap levels s.t. each circle maps to ~1 texel




Elliptical weighted averaging algorithm

1. approximate the ellipse with circles



2. find two appropriate mipmap levels s.t. each circle maps to ~1 texel




3. linearly interpolate between pixels and mipmap levels




How do we obtain the derivatives?

$$u, v = T(x_0, y_0) + \begin{bmatrix} \frac{\partial u}{\partial x} & \frac{\partial u}{\partial y} \\ \frac{\partial v}{\partial x} & \frac{\partial v}{\partial y} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}$$



The mapping T is the ray casting function

- so we can simply apply chain rule and differentiate ray casting



$$u, v = T(x_0, y_0) + \begin{bmatrix} \frac{\partial u}{\partial x} & \frac{\partial u}{\partial y} \\ \frac{\partial v}{\partial x} & \frac{\partial v}{\partial y} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}$$

The mapping T is the ray casting function

- so we can simply apply chain rule and differentiate ray casting

Tracing Ray Differentials

Homan Igehy

Computer Science Department

Stanford University





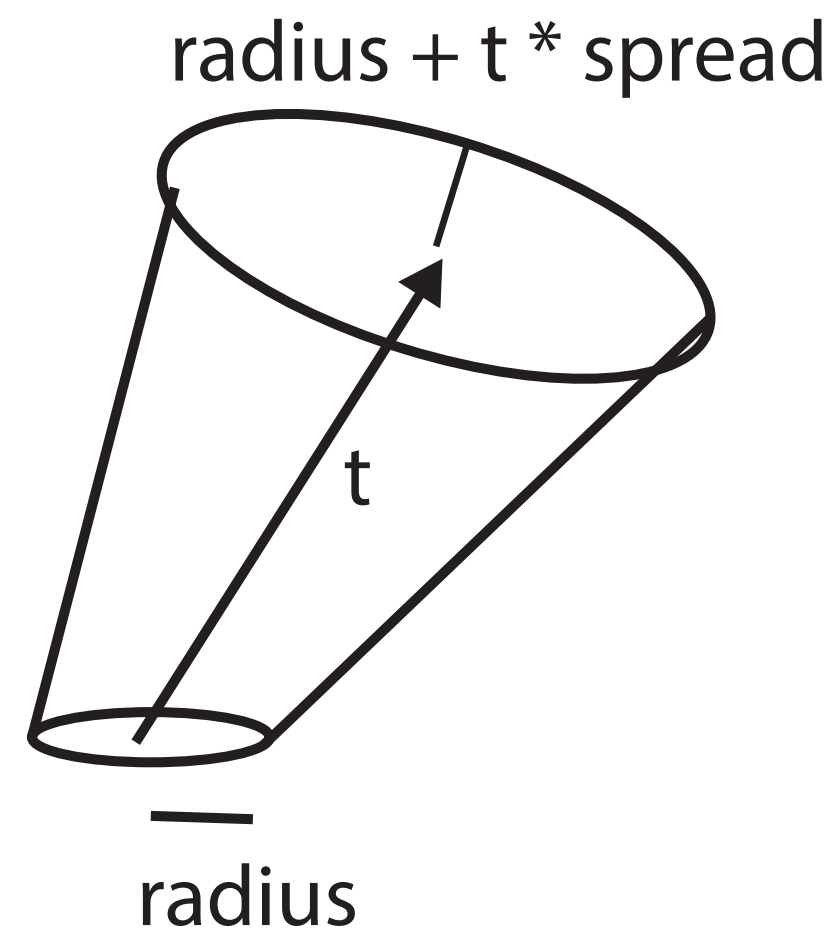
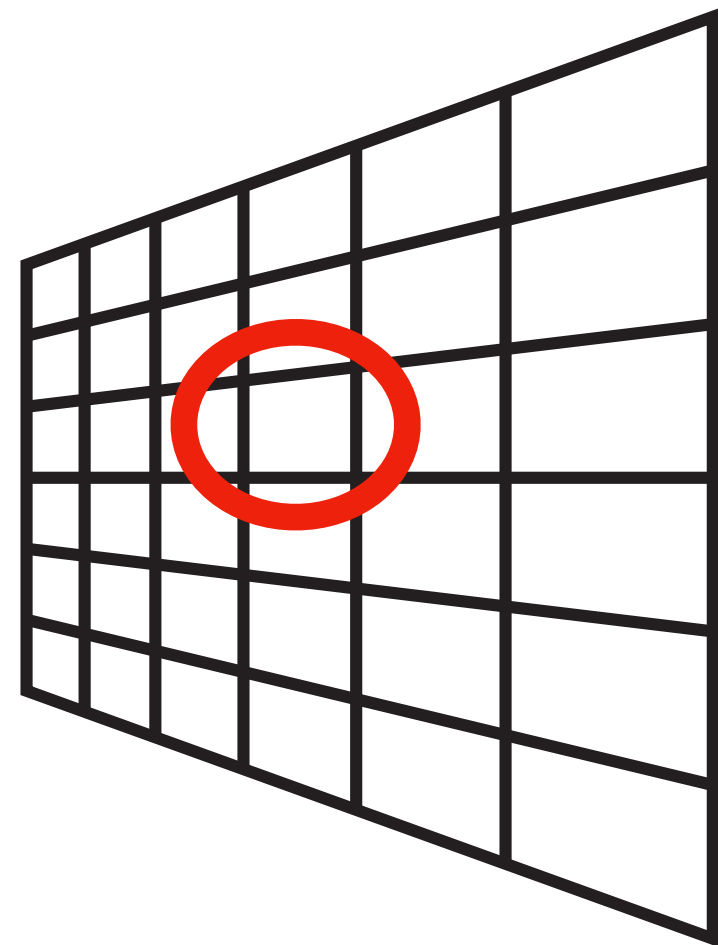
Figure 1: A Ray Differential. The diagram above illustrates the positions and directions of a ray and a differentially offset ray after a reflection. The difference between these positions and directions represents a ray differential.

$$u, v = T(x_0, y_0) + \begin{bmatrix} \frac{\partial u}{\partial x} & \frac{\partial u}{\partial y} \\ \frac{\partial v}{\partial x} & \frac{\partial v}{\partial y} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix}$$

Lajolla uses a heuristic method to approximate the ellipse with a circle

- see homework 0 for more details and the reasoning

initial spread = (pixel size) / 4



more about texture filtering next Monday!

https://github.com/BachiLi/lajolla_public/blob/main/src/ray.h
https://github.com/BachiLi/lajolla_public/blob/main/src/texture.h

Texture mapping: pros and cons

- pros
 - different sampling rates for geometry and color
 - much easier to filter
- cons
 - uv mapping is hard



In movie production: use UV generated by mesh subdivision

- will talk more about this in the later lectures




Ptex: Per-Face Texture Mapping for Production Rendering

Brent Burley^{†1} and Dylan Lacewell^{†1,2}

¹Walt Disney Animation Studios

²University of Utah




<https://www.youtube.com/watch?v=GxNlAlOuQQQ>

Figure 4: *a) Portion of a control mesh showing intrinsic faceids and edgeids. b) Corresponding limit surface showing continuous isolines across faces.*



Shading normals

- triangle meshes are planar approximations of a smooth surface: can look faceted.


quiz: how would you solve this?




Trick: assign a normal per vertex, then
interpolate



$$n = \text{normalize} \left((1 - b_1 - b_2)n_0 + b_1n_1 + b_2n_2 \right)$$



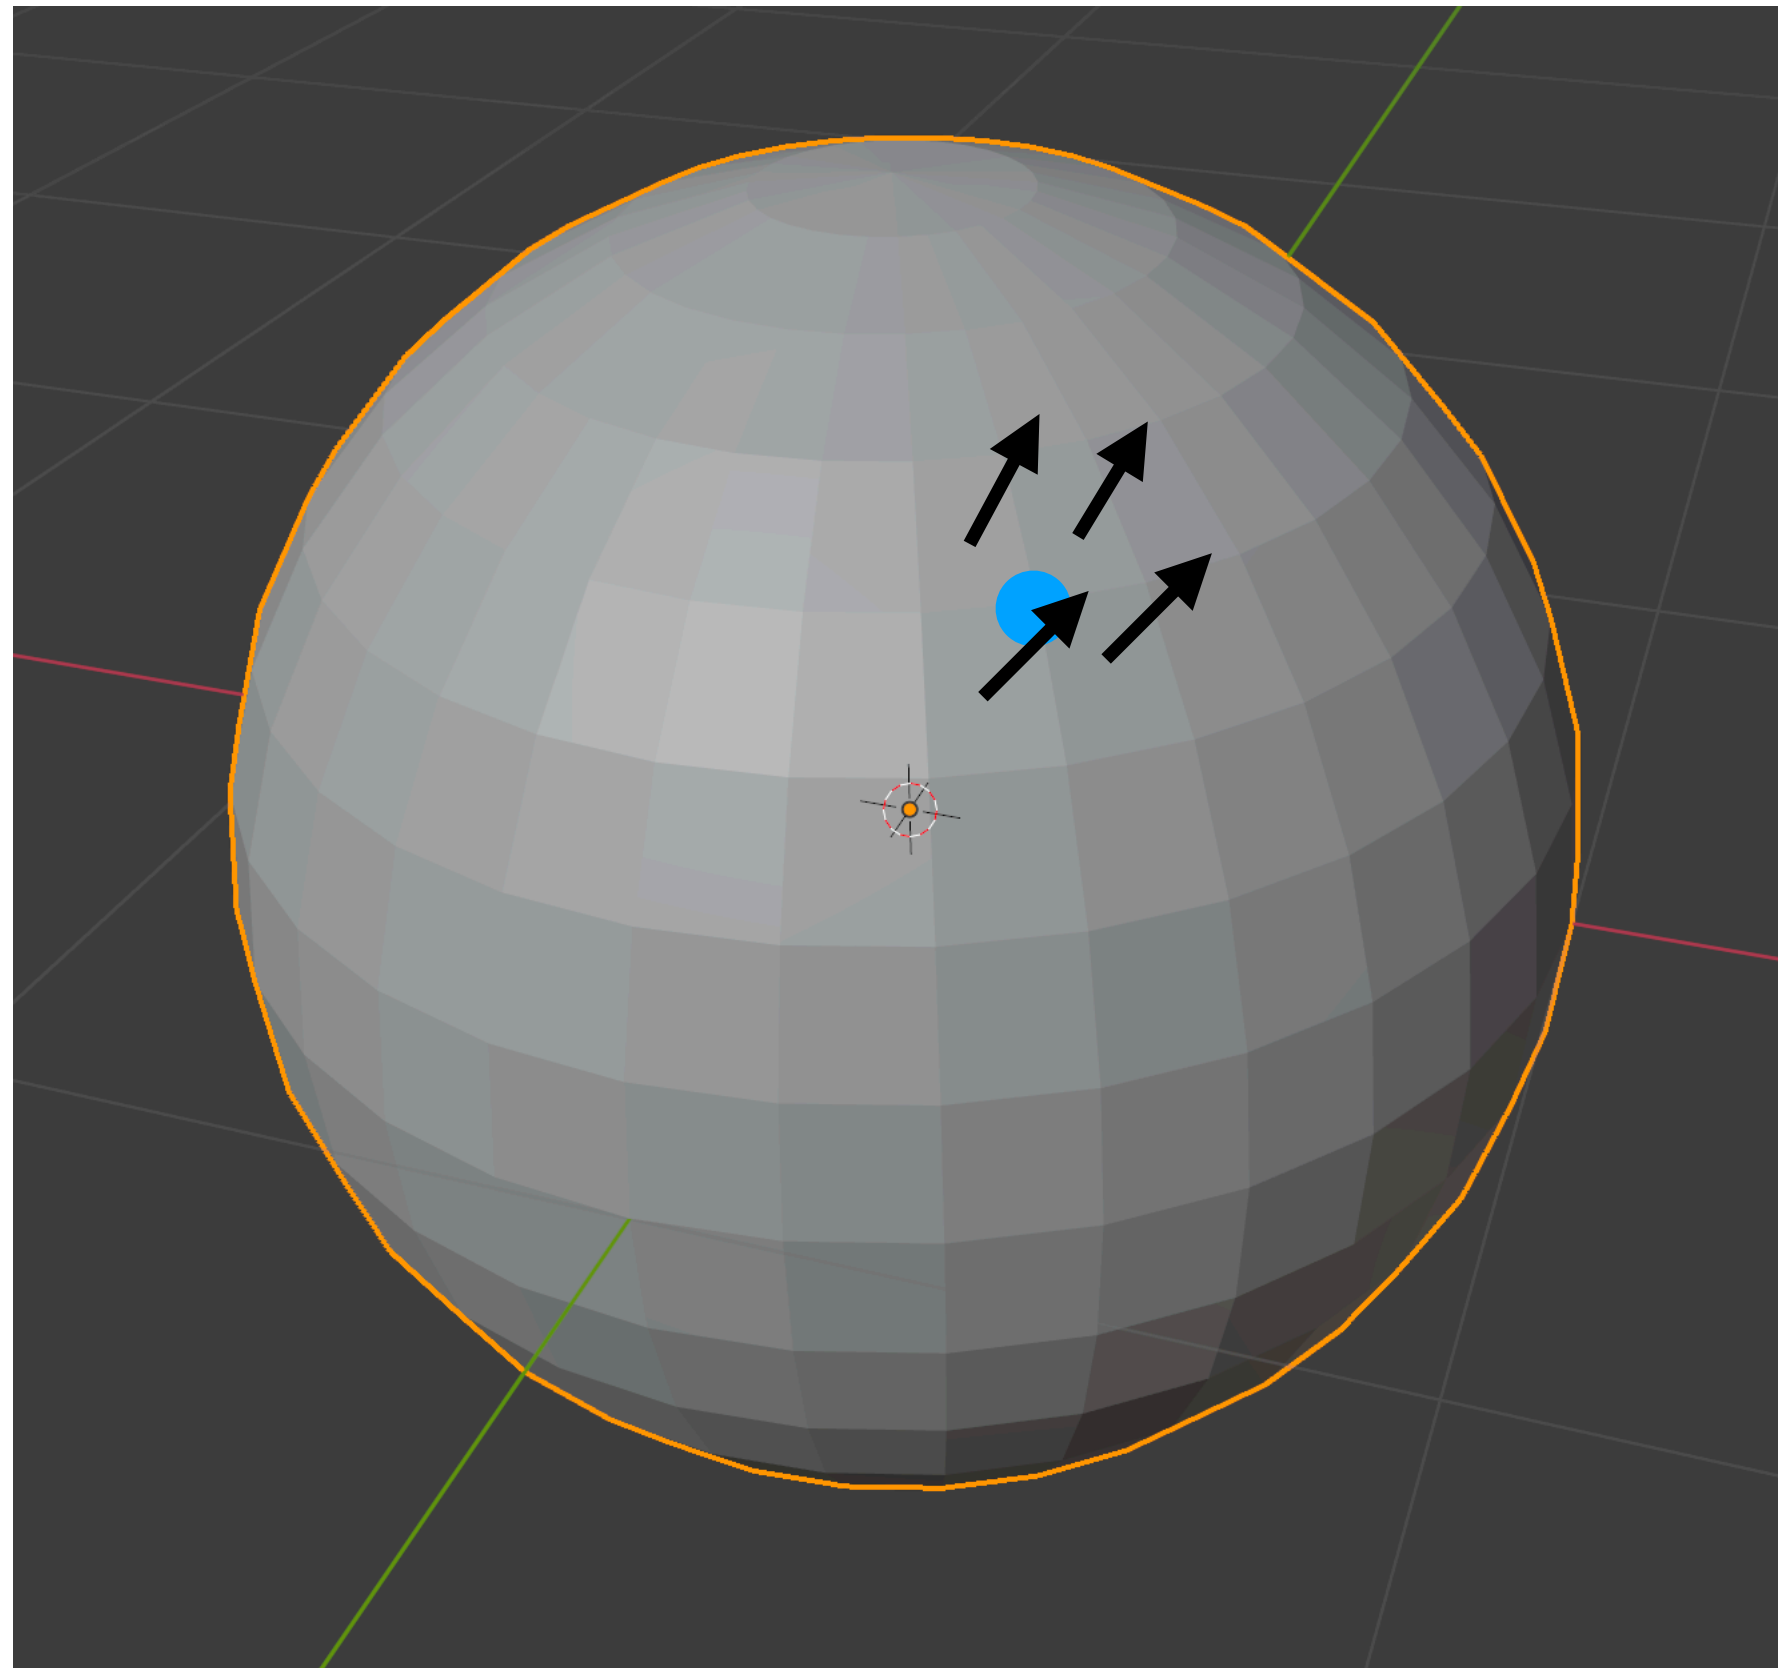
w / o per-vertex normal



w / per-vertex normal

How to get vertex normal?

- weighted average of normals of nearby faces





Weights for Computing Vertex Normals from Facet Normals


Nelson Max

Lawrence Livermore National Laboratory


Alternative: normal mapping



Discrepancies between shading normal and real geometry can lead to artifacts



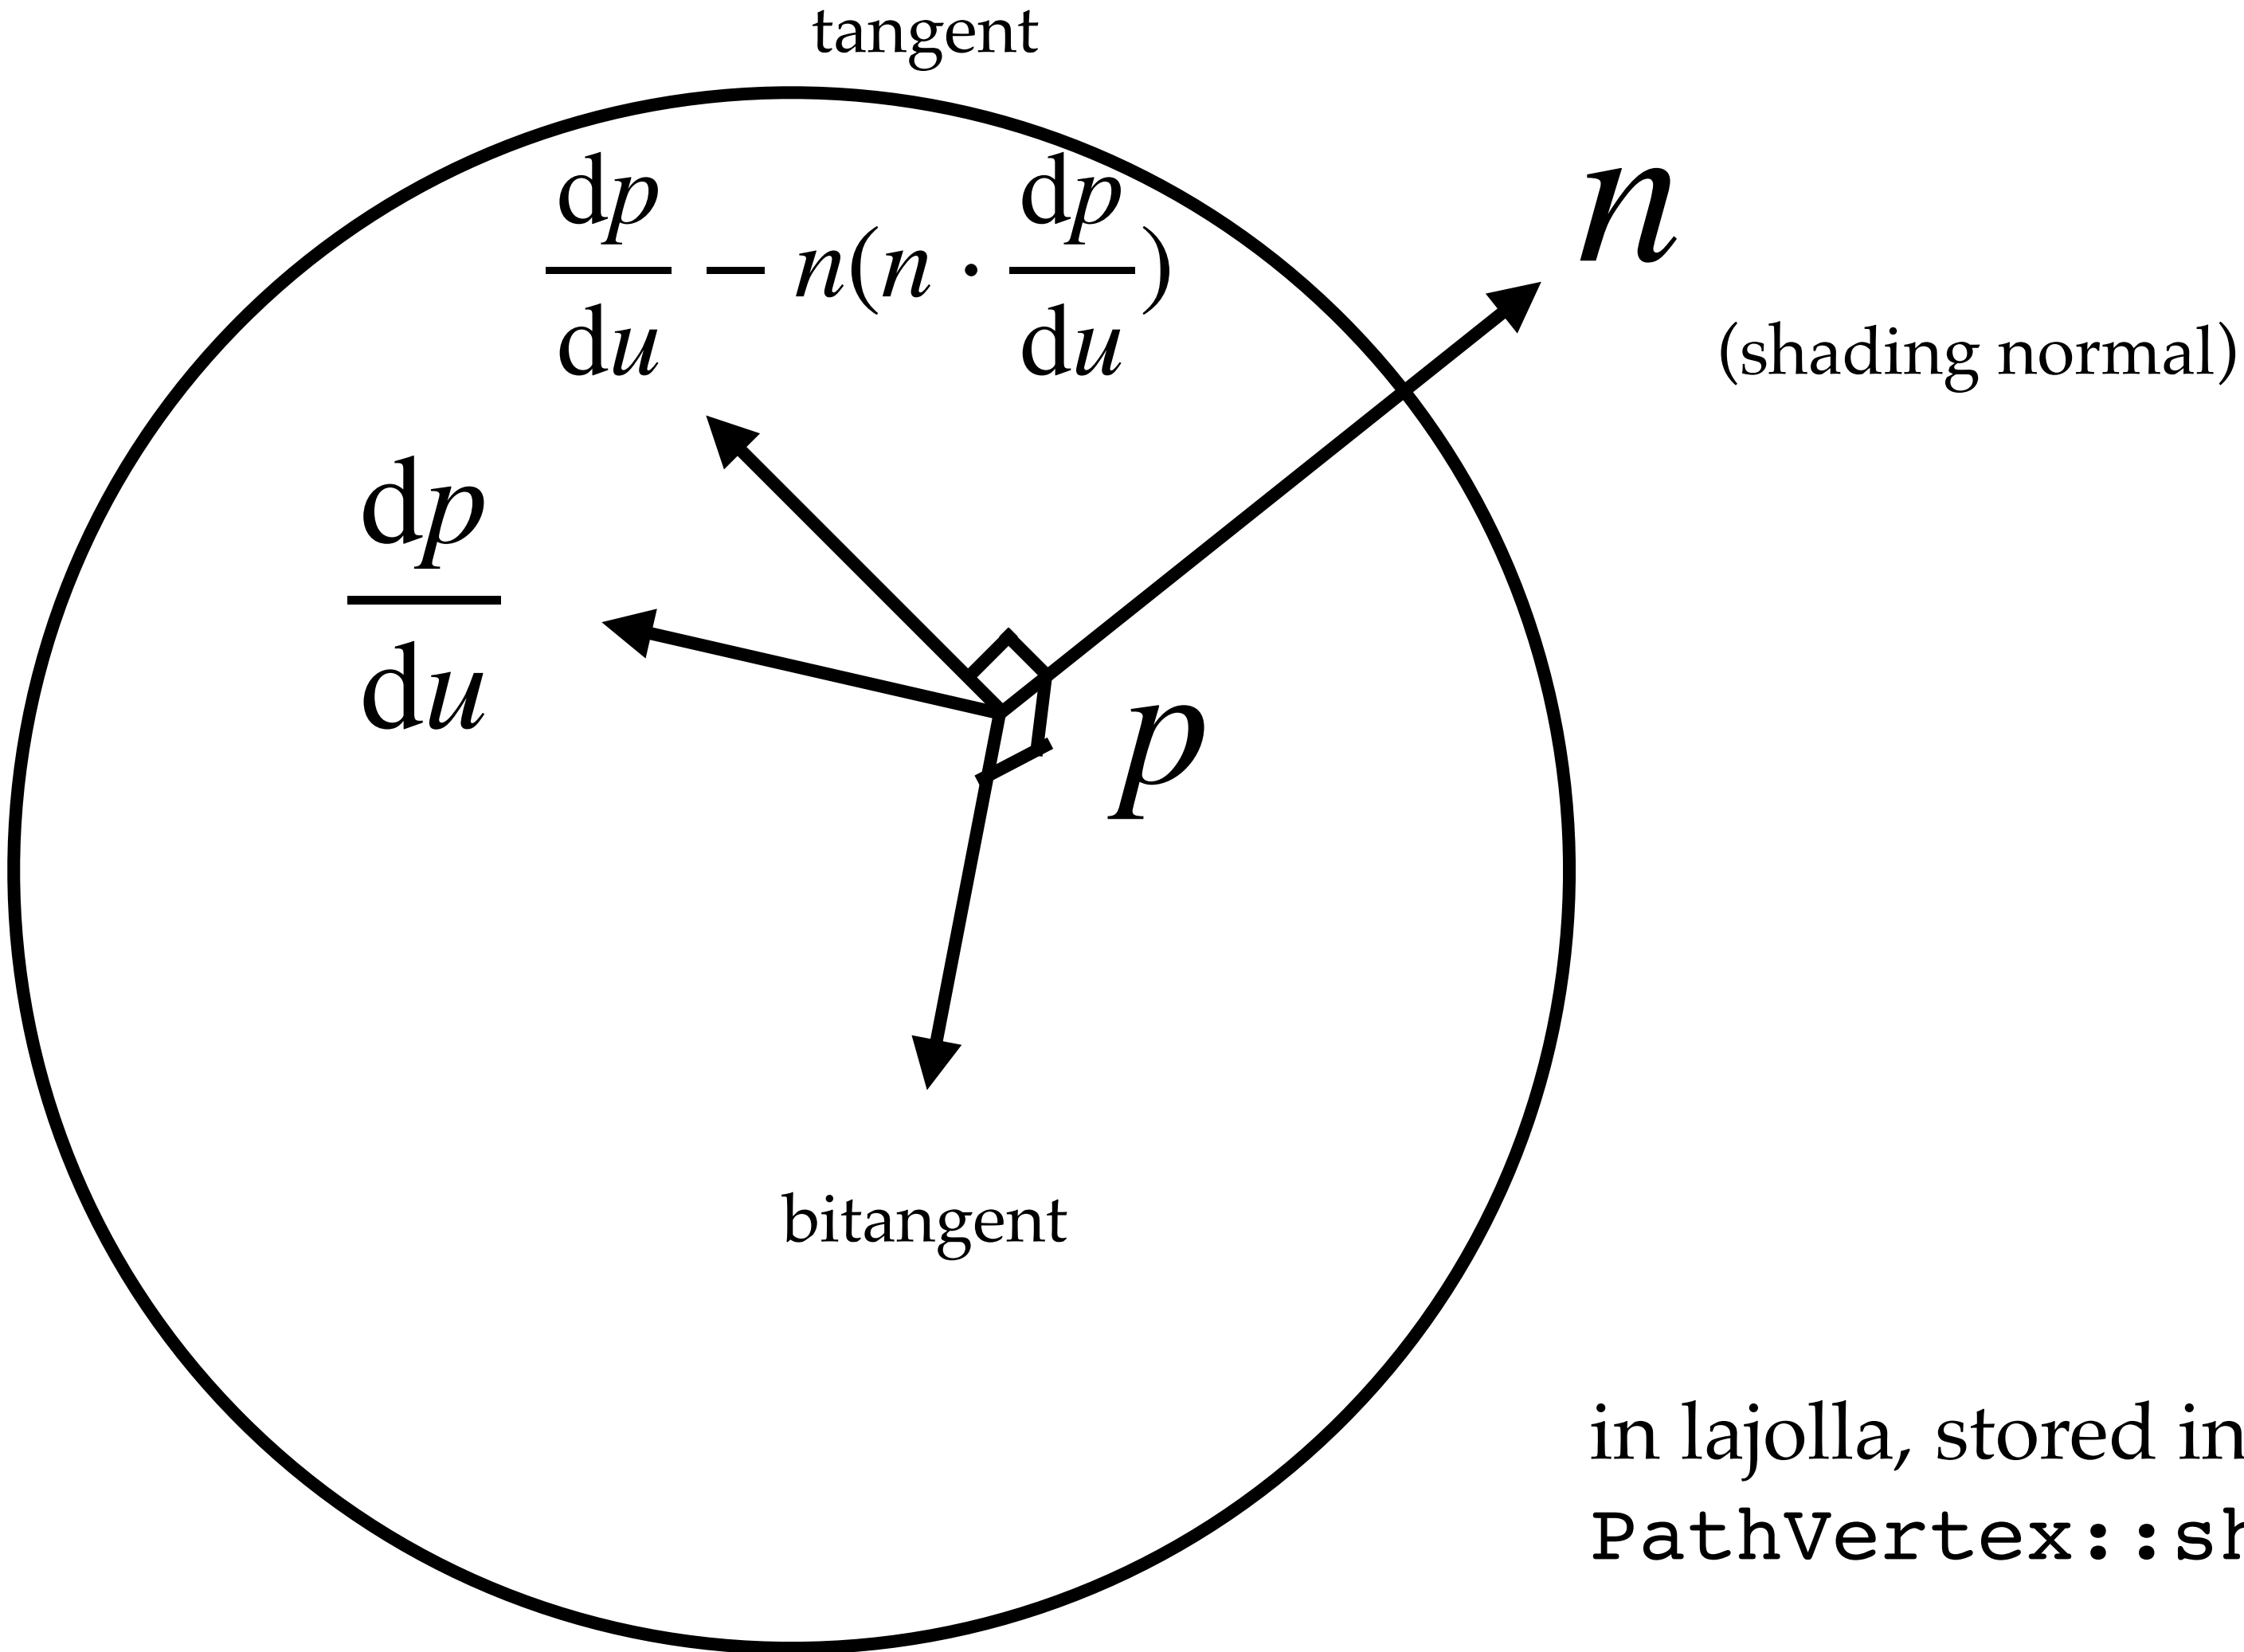
Discrepancies between shading normal and real geometry can lead to artifacts



Lajolla doesn't implement this!

From "Hacking the Shadow Terminator", Johannes Hanika
https://jo.dreggn.org/home/2021_terminator.pdf

For shading, we need a local coordinate basis



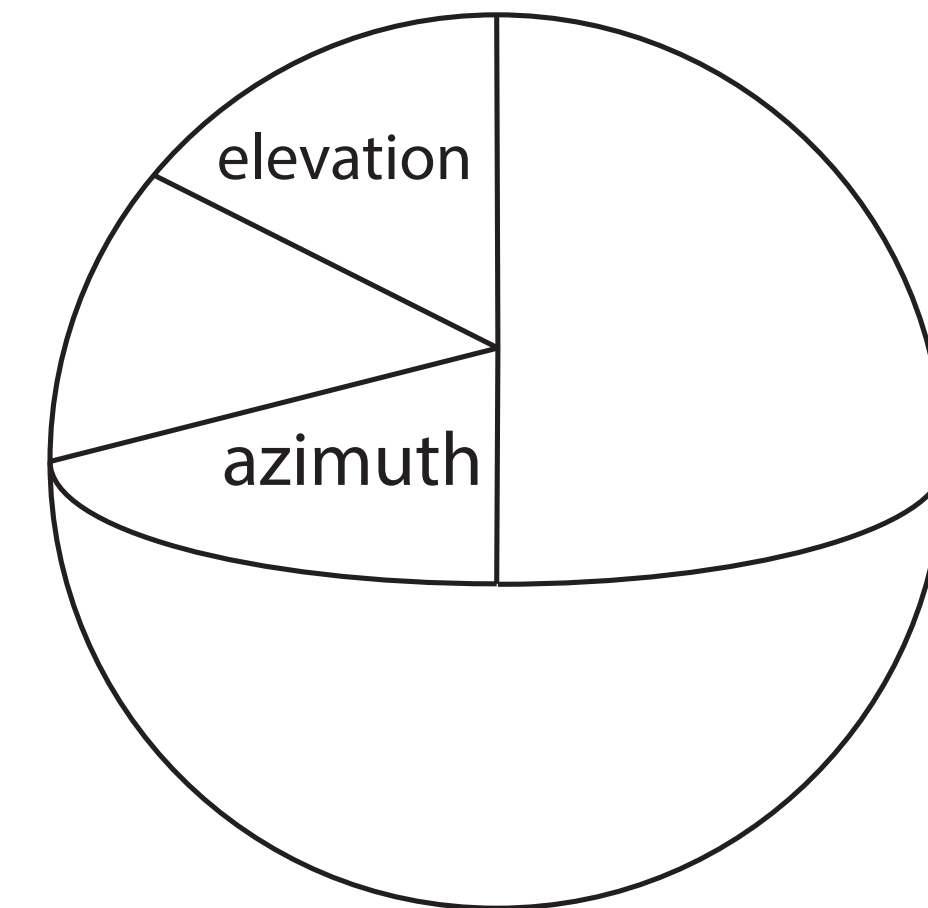
in lajolla, stored in
`PathVertex::shading_frame`

In Iajolla

```
struct PathVertex {  
    // ...  
    Vector3 geometry_normal; // always face at the same direction at shading_frame.n  
    Frame shading_frame;  
    // ...  
};
```


Environment maps

- an infinitely far area light source represented as an image




Environment maps

- an infinitely far area light source represented as an image



Environment maps sampling: treat it as a big discrete 2D table

- first sample a row, then sample a column



Blender-Mitsuba converter

mitsuba-renderer / mitsuba-blender Public Watch 14

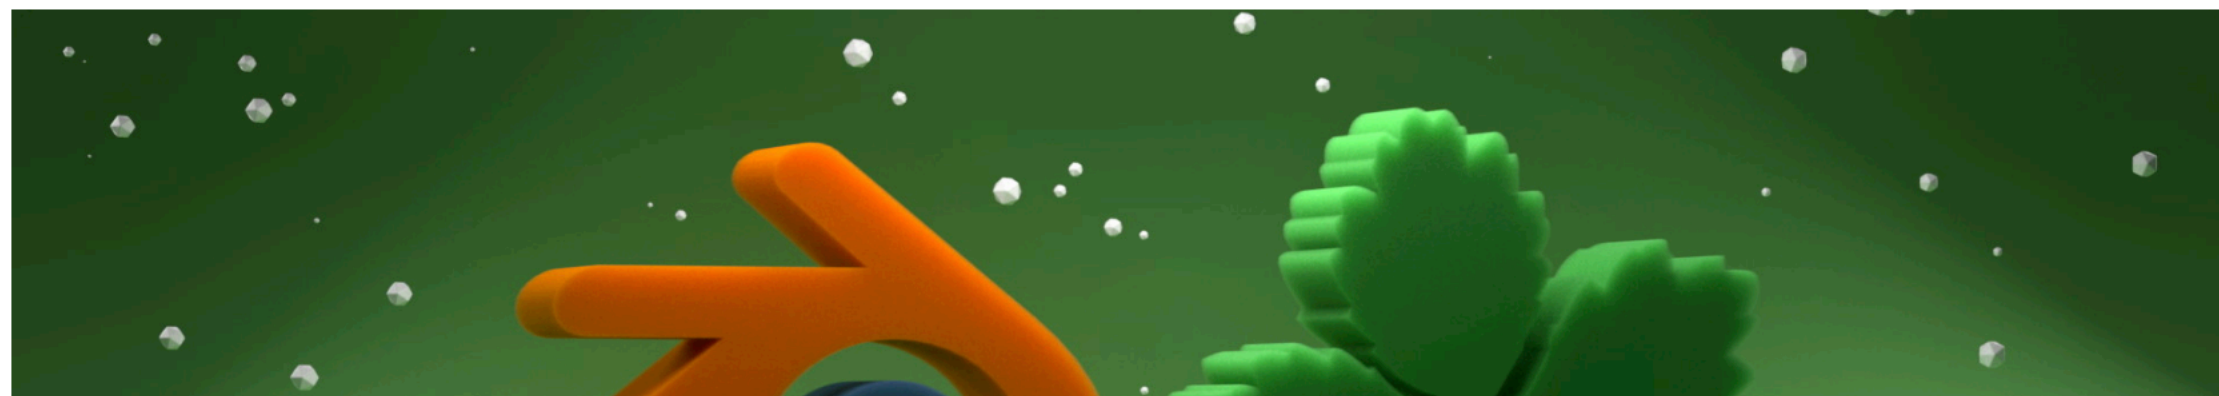
[Code](#) [Issues 8](#) [Pull requests 2](#) [Discussions](#) [Actions](#) [Projects](#) [Wiki](#) [Security](#) [Insights](#)

master 2 branches 2 tags Go to file Add file Code

ros-dorian Importer: Fix y-axis UV flip with the PLY object importer (#51) ... ✓ 5985305 on Nov 22, 2022 🕒 290 commits

📁 .github	Installation: Define a fixed version of Mitsuba dependencies (#44)	3 months ago
📁 mitsuba-blender	Importer: Fix y-axis UV flip with the PLY object importer (#51)	last month
📁 release	Add nightly release workflow	4 months ago
📁 res	Meta: Simplify installation instructions in README	4 months ago
📁 scripts	Rename occurrences of Mitsuba 2 to Mitsuba	4 months ago
📁 tests	Rename occurrences of Mitsuba 2 to Mitsuba	4 months ago
📄 .gitignore	Add support for custom Pytest arguments	8 months ago
📄 LICENSE	Added LICENSE	8 months ago
📄 README.md	Meta: Switch status badge to nightly release workflow	3 months ago

☰ README.md



The screenshot shows a 3D rendered scene with a green background, white particles, and a large orange hand-like object reaching towards a green plant-like structure.

About

Mitsuba integration add-on for Blender


[mitsuba](#) [blender-addon](#)

📖 Readme
📄 BSD-3-Clause license
★ 153 stars
👁 14 watching
🍴 19 forks

Releases

🏷 2 tags

Contributors 5





Five circular avatars of contributors are shown in a row.

Languages

Python 100.0%

Next time:

bidirectional scattering distribution function



$$L(\mathbf{p}, \omega) = L_e(\mathbf{p}, \omega) + \int \boxed{f_{\mathbf{p}}(\omega, \omega')} L(\mathbf{p}', -\omega') |n_{\mathbf{p}} \cdot \omega'| d\omega'$$