

Lecture 4:

Perspective Projection and Texture Mapping

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

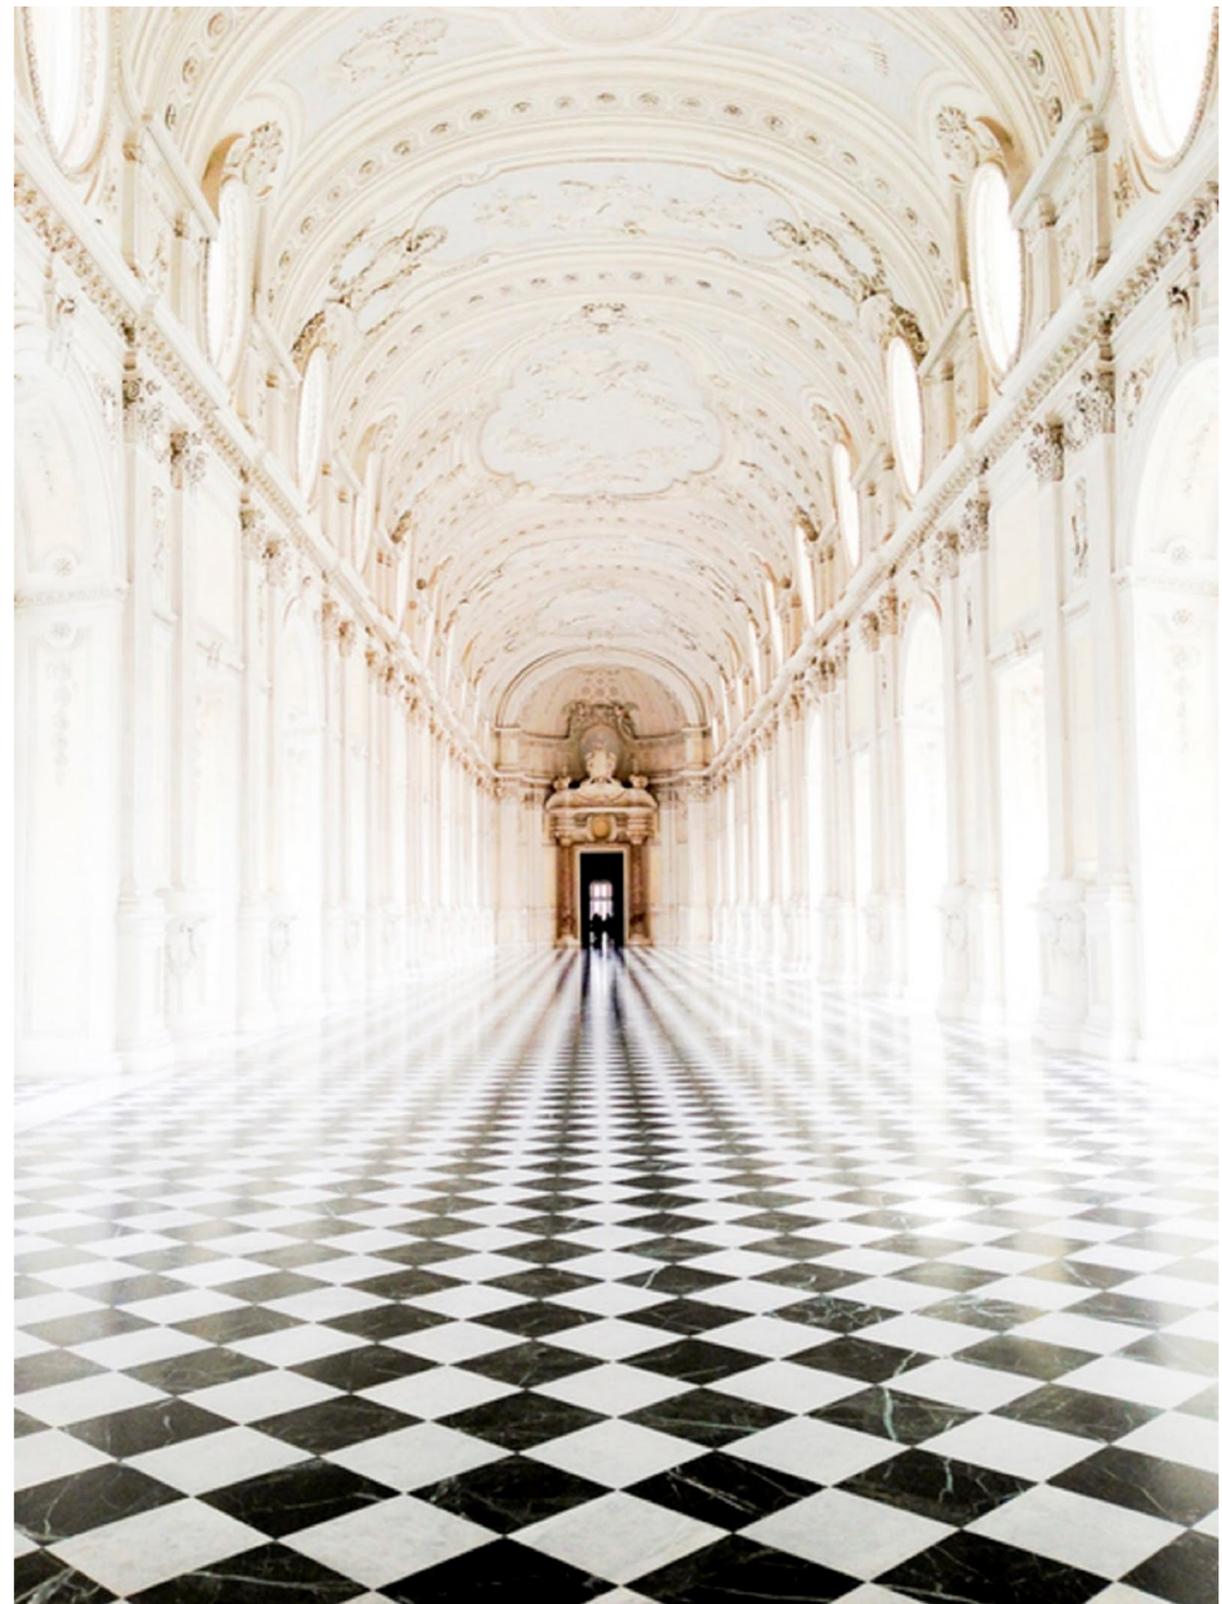
Perspective and texture

■ PREVIOUSLY:

- *transformation* (how to manipulate primitives in space)
- *rasterization* (how to turn primitives into pixels)

■ TODAY:

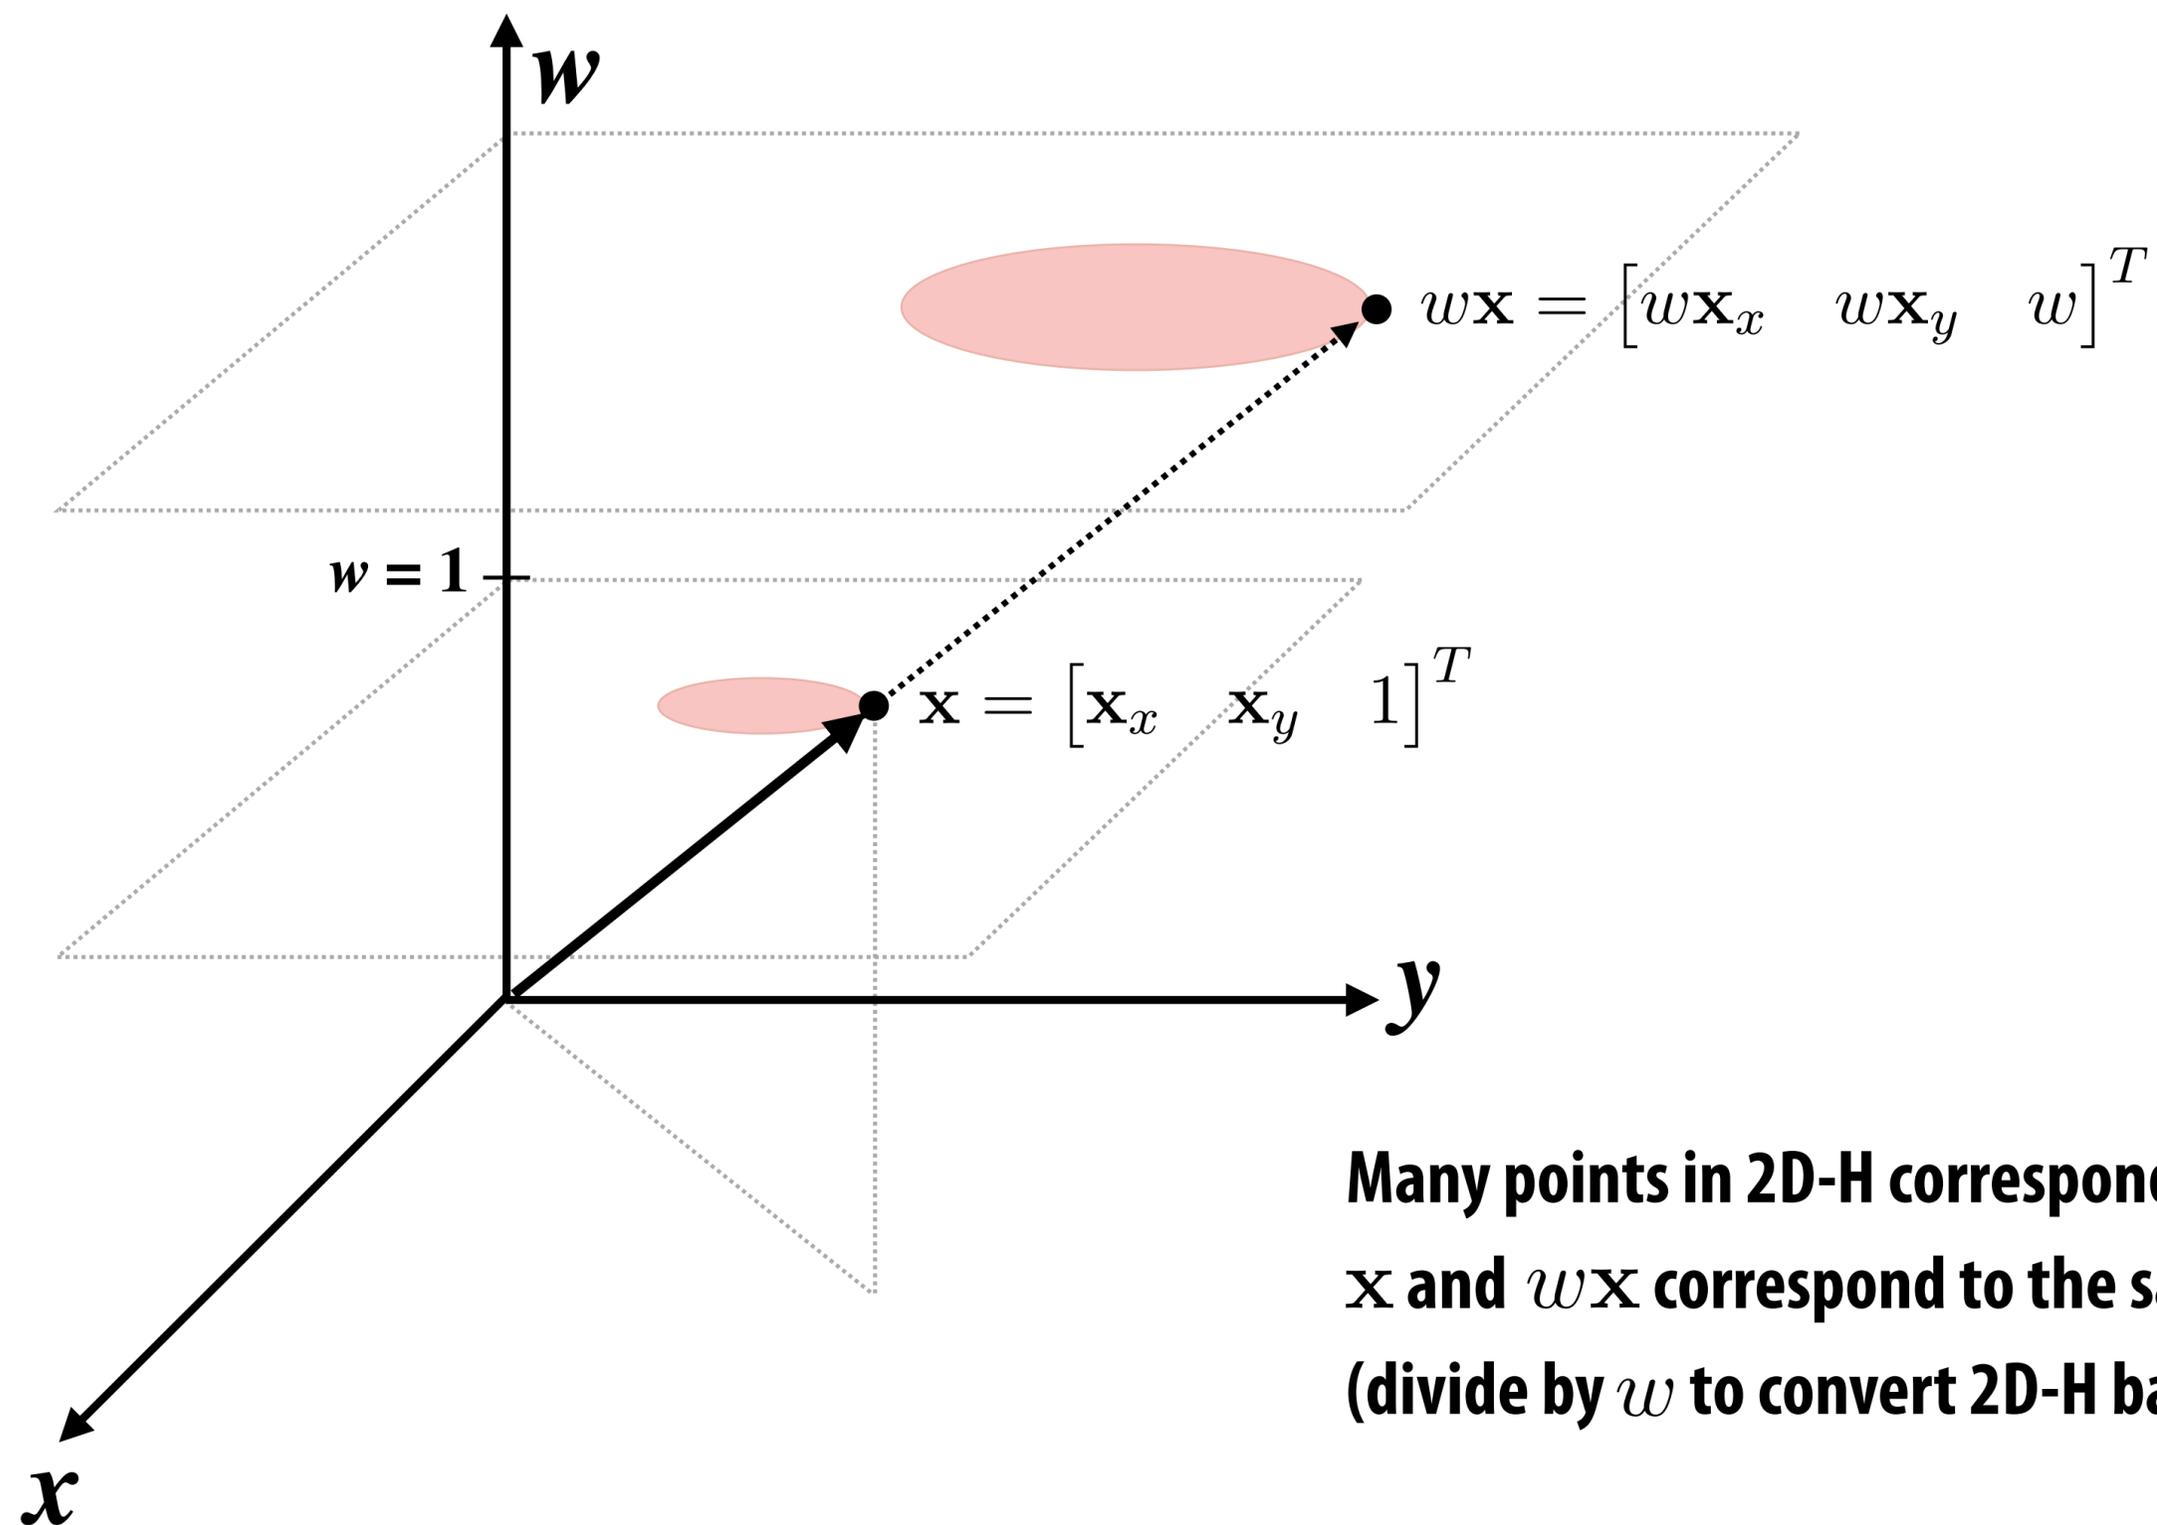
- see where these two ideas come crashing together!
- revisit *perspective* transformations
- talk about how to map *texture* onto a primitive to get more detail
- ...and how perspective creates challenges for texture mapping!



**Why is it hard to render
an image like this?**

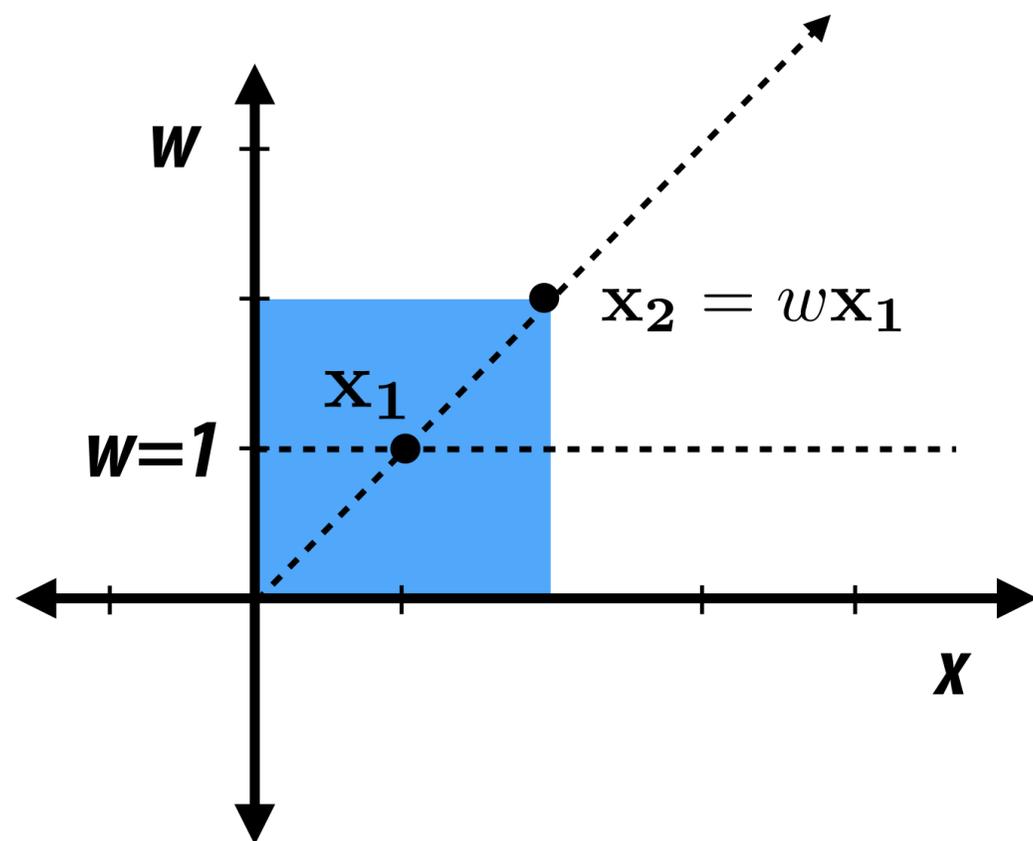
Transformations review

Review: homogeneous coordinates



**Many points in 2D-H correspond to same point in 2D:
 \mathbf{x} and $w\mathbf{x}$ correspond to the same 2D point
(divide by w to convert 2D-H back to 2D)**

Translation = shear in homogeneous space

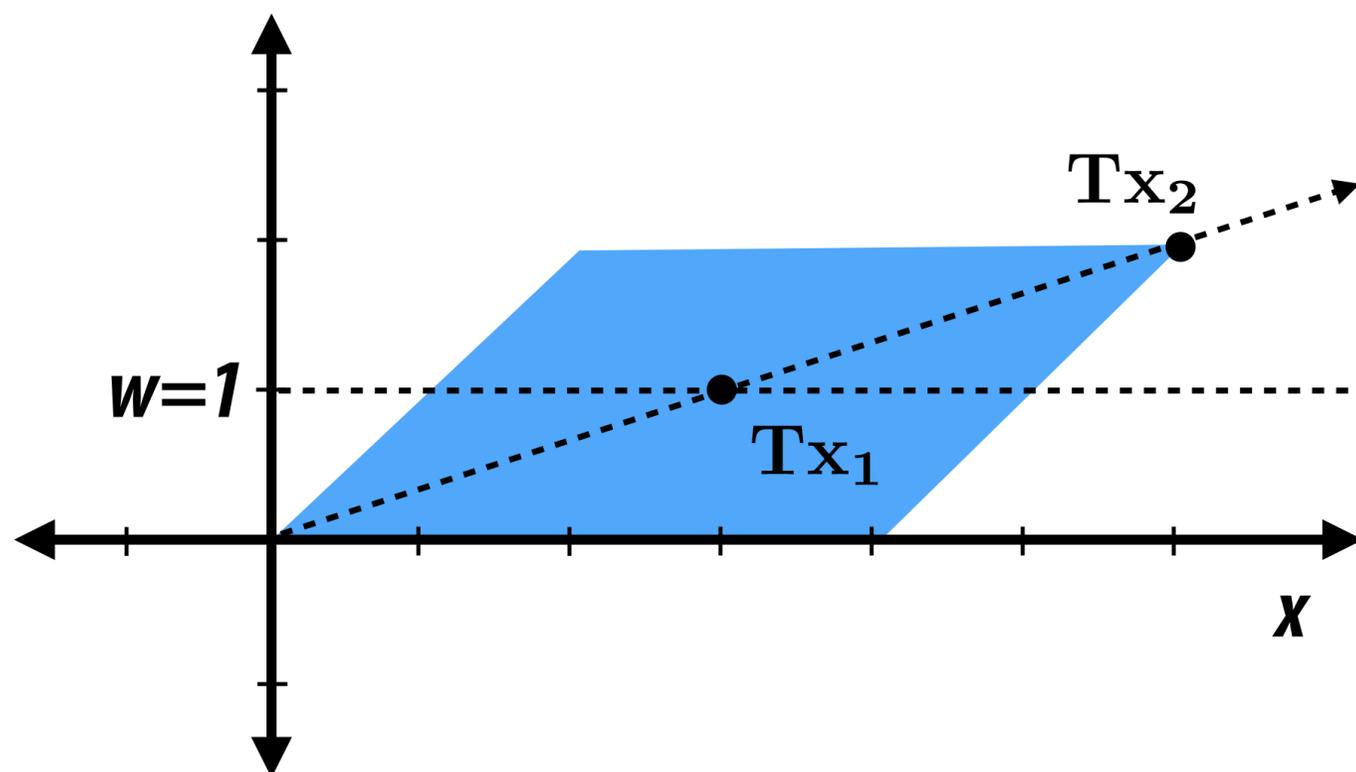


Consider 1D-H:

$$\text{Translate by } t=2: \mathbf{T} = \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} 1 & 2 \\ 0 & 1 \end{bmatrix}$$

Recall: this is a shear in homogeneous x .

1D translation is affine in 1D ($x + t$),
but it is linear in 1D-H



Skeleton - hierarchical representation

torso

head

right arm

upper arm

lower arm

hand

left arm

upper arm

lower arm

hand

right leg

upper leg

lower leg

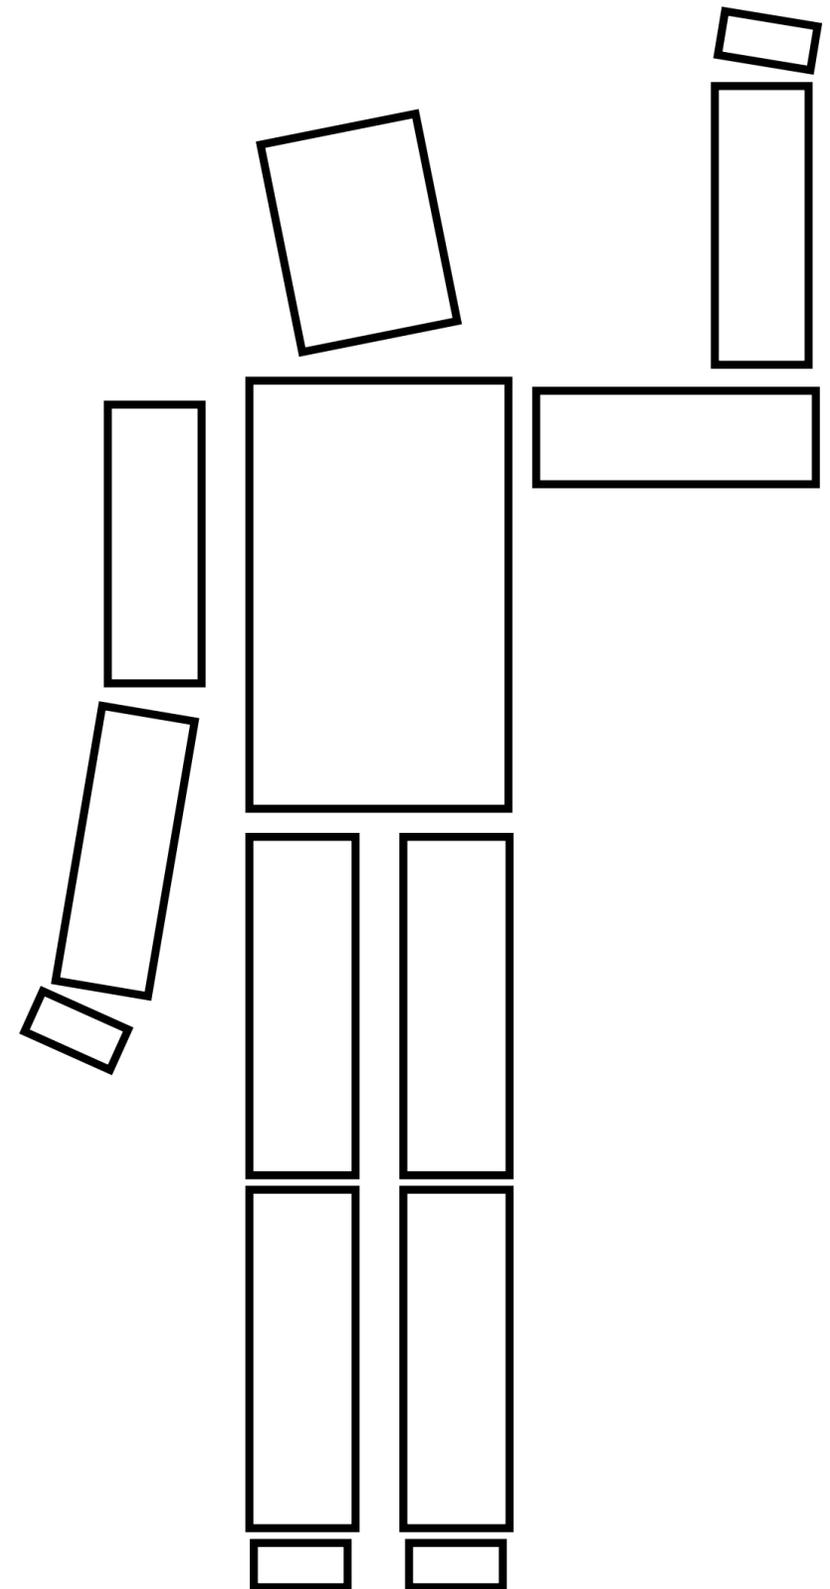
foot

left leg

upper leg

lower leg

foot



Hierarchical representation

- **Grouped representation (tree)**
 - **Each group contains subgroups and/or shapes**
 - **Each group is associated with a transform relative to parent group**
 - **Transform on leaf-node shape is concatenation of all transforms on path from root node to leaf**
 - **Changing a group's transform affects all parts**
 - **Allows high level editing by changing only one node**
 - **E.g. raising left arm requires changing only one transform for that group**

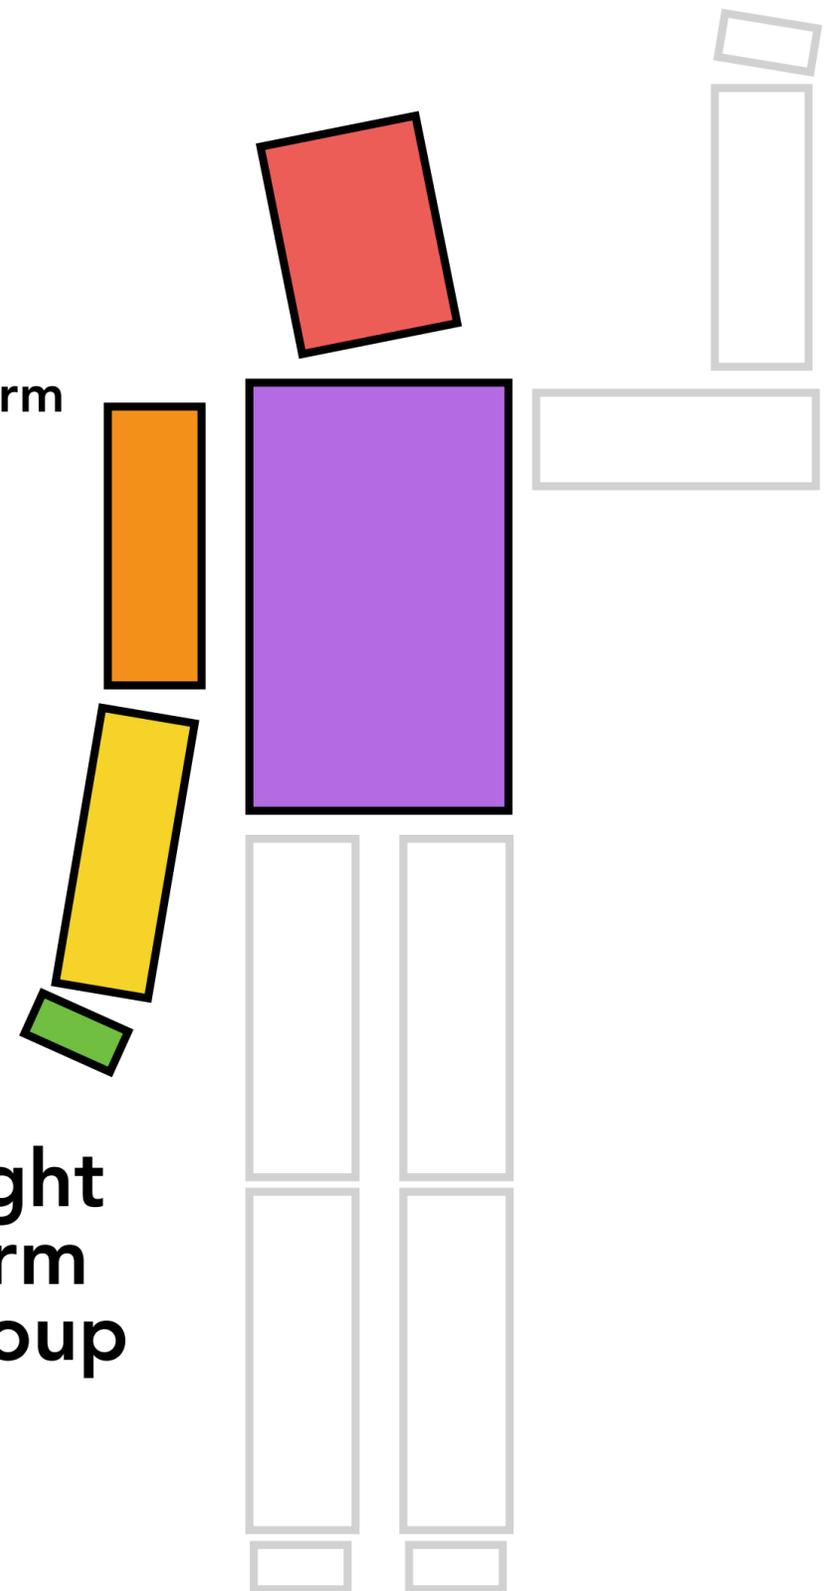
Skeleton - hierarchical representation

```
translate(0, 10);
drawTorso();
pushmatrix(); // push a copy of transform onto stack
  translate(0, 5); // right-multiply onto current transform
  rotate(headRotation); // right-multiply onto current transform
  drawHead();
popmatrix(); // pop current transform off stack
pushmatrix();
  translate(-2, 3);
  rotate(rightShoulderRotation);
  drawUpperArm();
  pushmatrix();
    translate(0, -3);
    rotate(elbowRotation);
    drawLowerArm();
    pushmatrix();
      translate(0, -3);
      rotate(wristRotation);
      drawHand();
    popmatrix();
  popmatrix();
popmatrix();
....
```

right
hand

right
lower
arm
group

right
arm
group



Skeleton - hierarchical representation

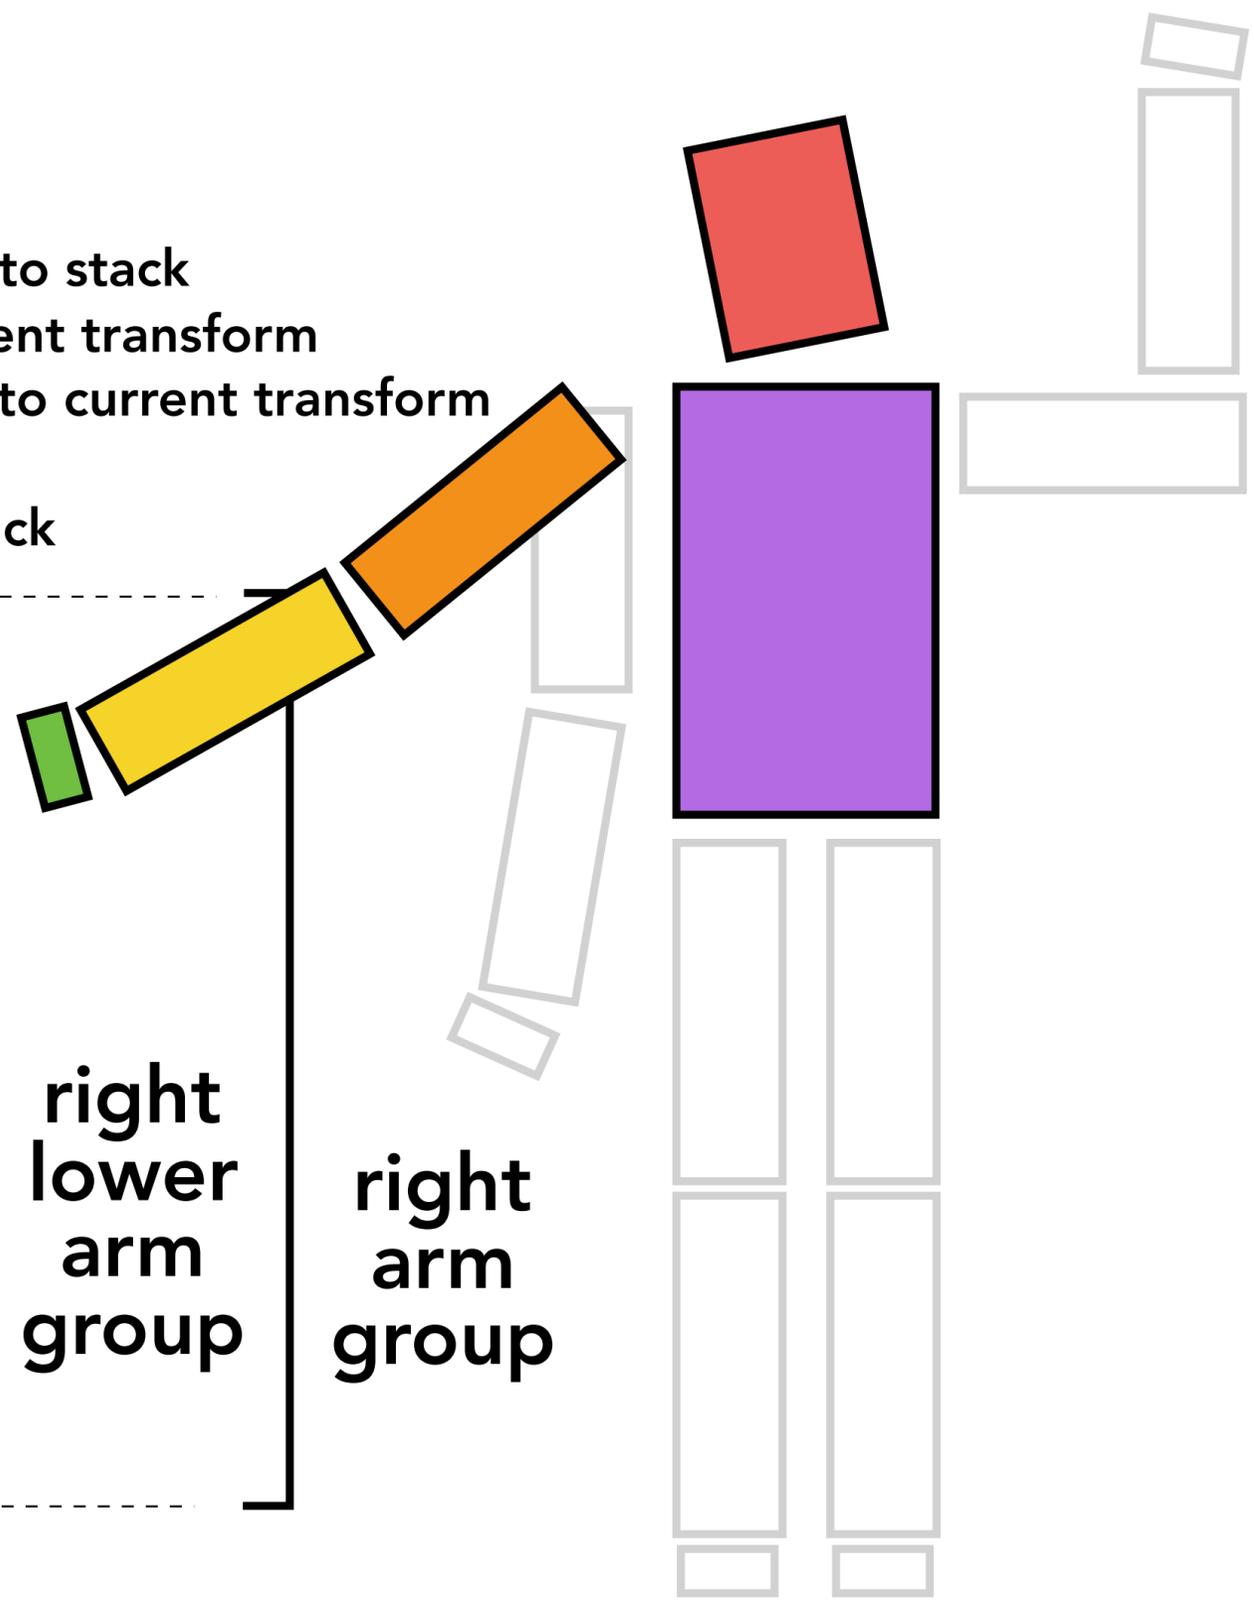
```
translate(0, 10);
drawTorso();
pushmatrix(); // push a copy of transform onto stack
  translate(0, 5); // right-multiply onto current transform
  rotate(headRotation); // right-multiply onto current transform
  drawHead();
popmatrix(); // pop current transform off stack
pushmatrix();
  translate(-2, 3);
  rotate(rightShoulderRotation);
  drawUpperArm();
  pushmatrix();
    translate(0, -3);
    rotate(elbowRotation);
    drawLowerArm();
    pushmatrix();
      translate(0, -3);
      rotate(wristRotation);
      drawHand();
    popmatrix();
  popmatrix();
popmatrix();
....
```



right
hand

right
lower
arm
group

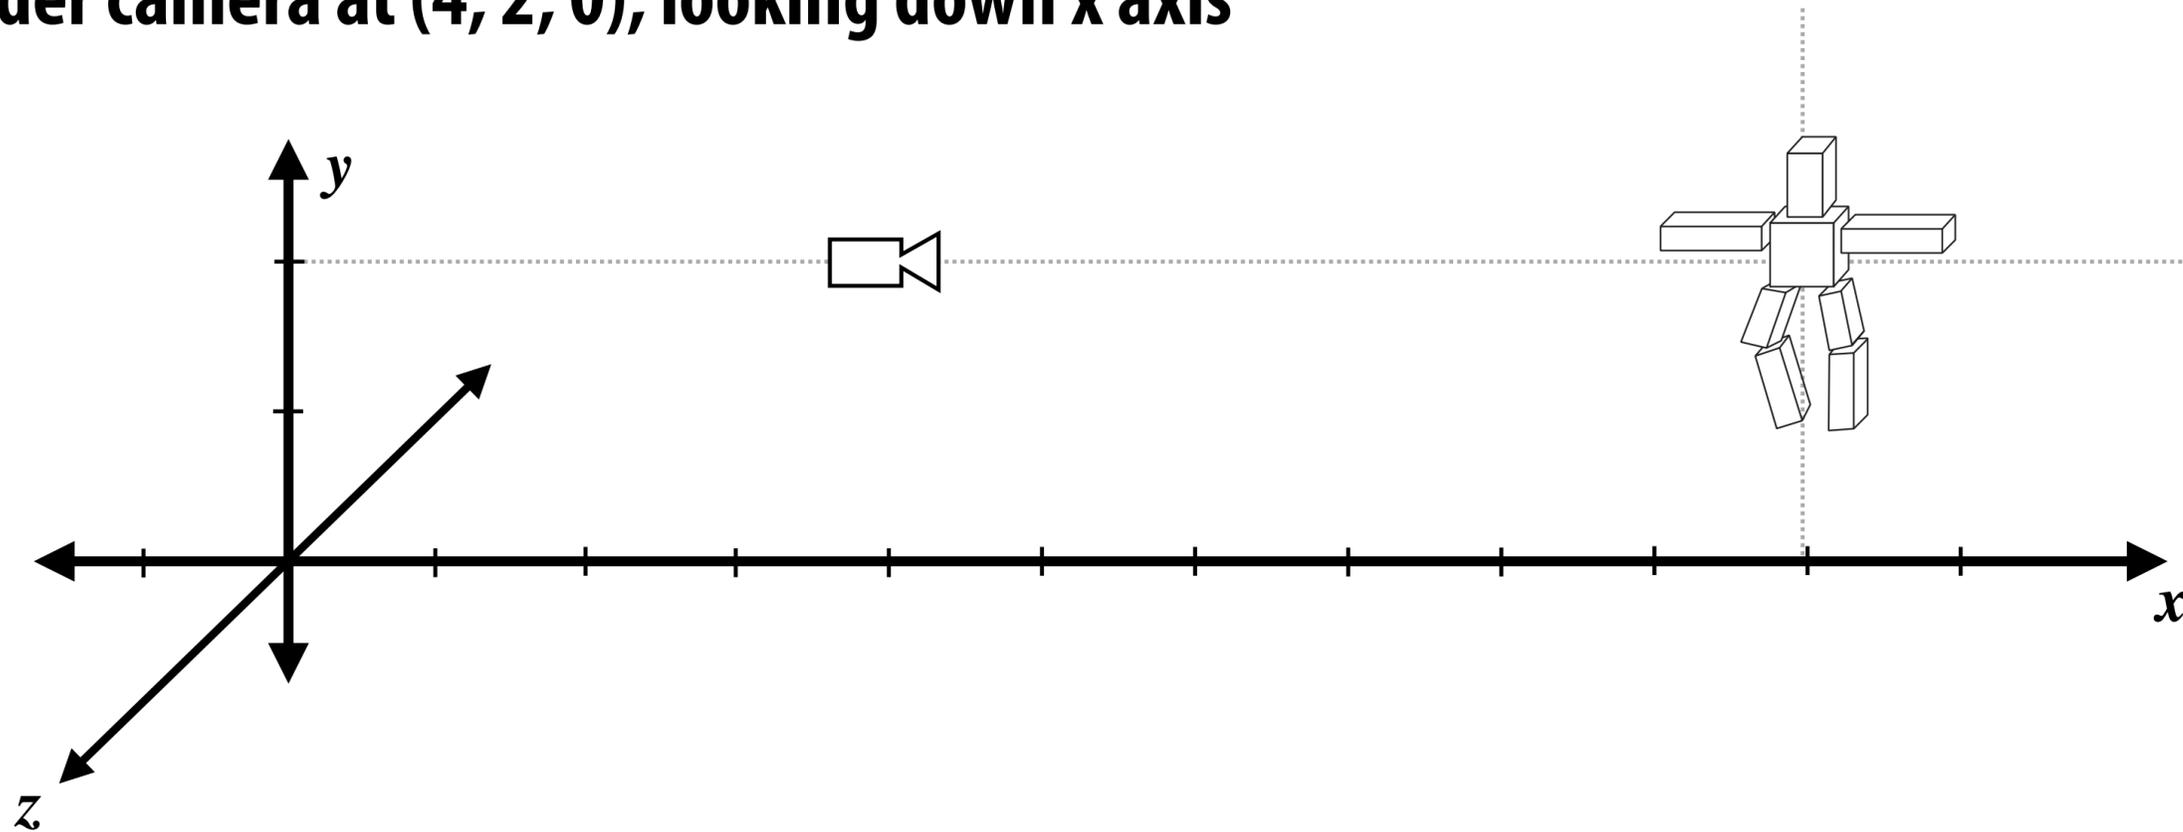
right
arm
group



Review: simple camera transform

Consider object positioned in world at $(10, 2, 0)$

Consider camera at $(4, 2, 0)$, looking down x axis



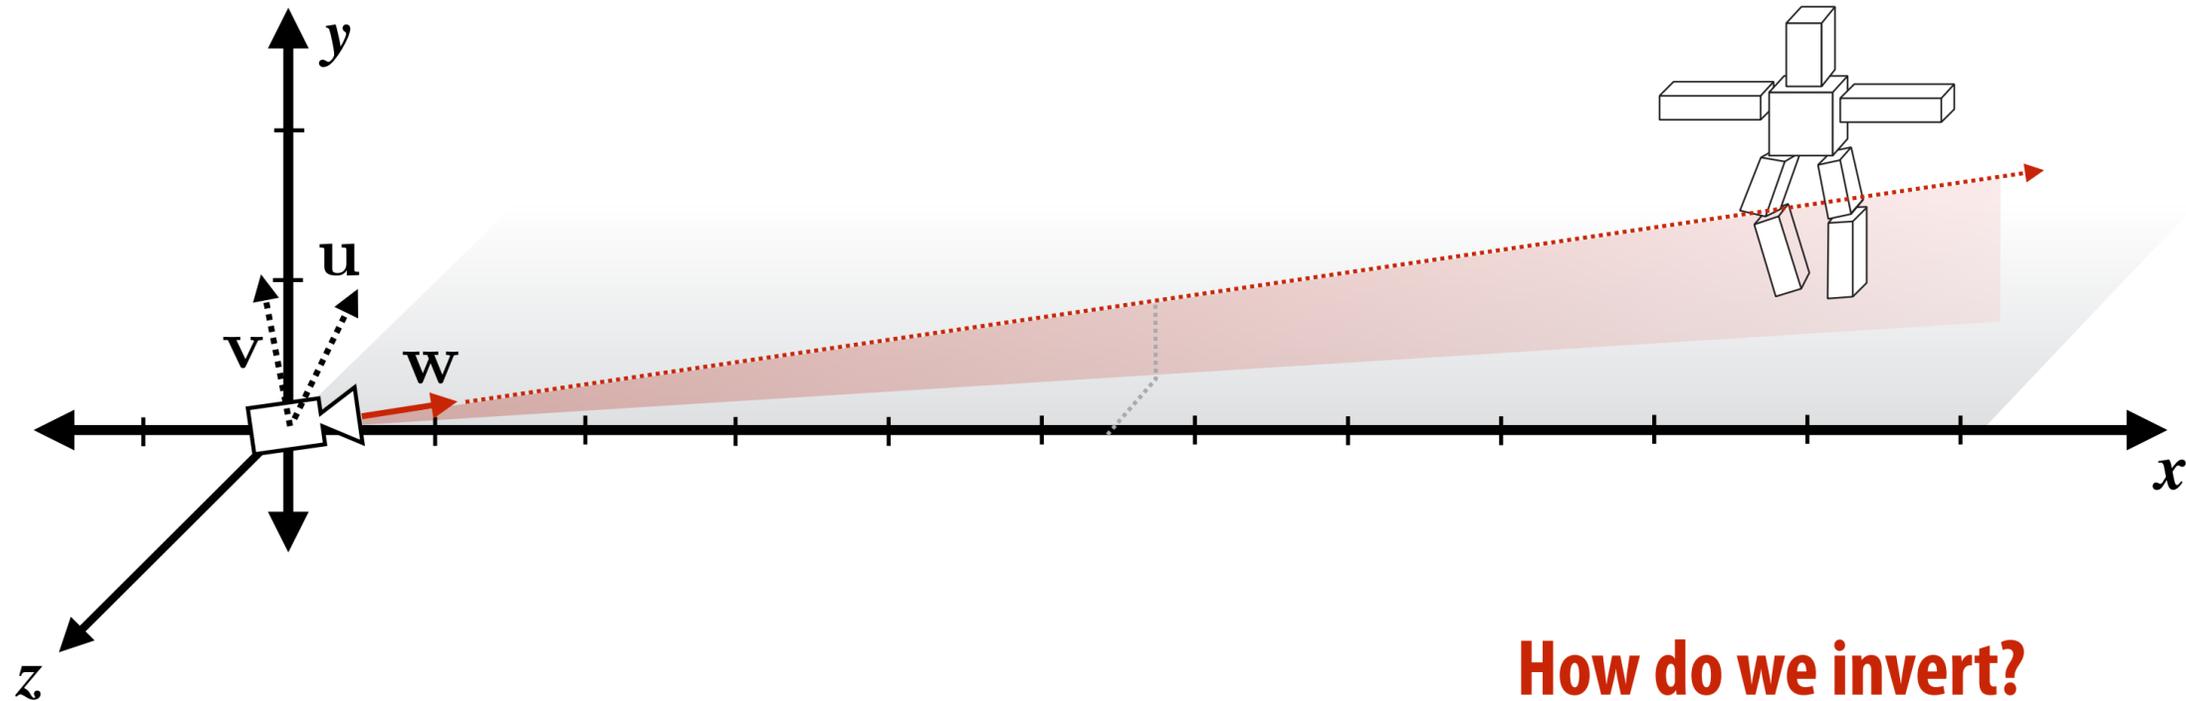
What transform places in the object in a coordinate space where the camera is at the origin and the camera is looking directly down the -z axis?

- Translating object vertex positions by $(-4, -2, 0)$ yields position relative to camera
- Rotation about y by $\pi/2$ gives position of object in new coordinate system where camera's view direction is aligned with the $-z$ axis

Camera looking in a different direction

Consider camera looking in direction \mathbf{w}

What transform places in the object in a coordinate space where the camera is at the origin and the camera is looking directly down the $-z$ axis?



Form orthonormal basis around \mathbf{w} : (see \mathbf{u} and \mathbf{v})

Consider rotation matrix: \mathbf{R}

$$\mathbf{R} = \begin{bmatrix} \mathbf{u}_x & \mathbf{v}_x & -\mathbf{w}_x \\ \mathbf{u}_y & \mathbf{v}_y & -\mathbf{w}_y \\ \mathbf{u}_z & \mathbf{v}_z & -\mathbf{w}_z \end{bmatrix}$$

\mathbf{R} maps x -axis to \mathbf{u} , y -axis to \mathbf{v} , z axis to $-\mathbf{w}$

How do we invert?

$$\mathbf{R}^{-1} = \mathbf{R}^T = \begin{bmatrix} \mathbf{u}_x & \mathbf{u}_y & \mathbf{u}_z \\ \mathbf{v}_x & \mathbf{v}_y & \mathbf{v}_z \\ -\mathbf{w}_x & -\mathbf{w}_y & -\mathbf{w}_z \end{bmatrix}$$

Why is that the inverse?

$$\mathbf{R}^T \mathbf{u} = [\mathbf{u} \cdot \mathbf{u} \quad \mathbf{v} \cdot \mathbf{u} \quad -\mathbf{w} \cdot \mathbf{u}]^T = [1 \quad 0 \quad 0]^T$$

$$\mathbf{R}^T \mathbf{v} = [\mathbf{u} \cdot \mathbf{v} \quad \mathbf{v} \cdot \mathbf{v} \quad -\mathbf{w} \cdot \mathbf{v}]^T = [0 \quad 1 \quad 0]^T$$

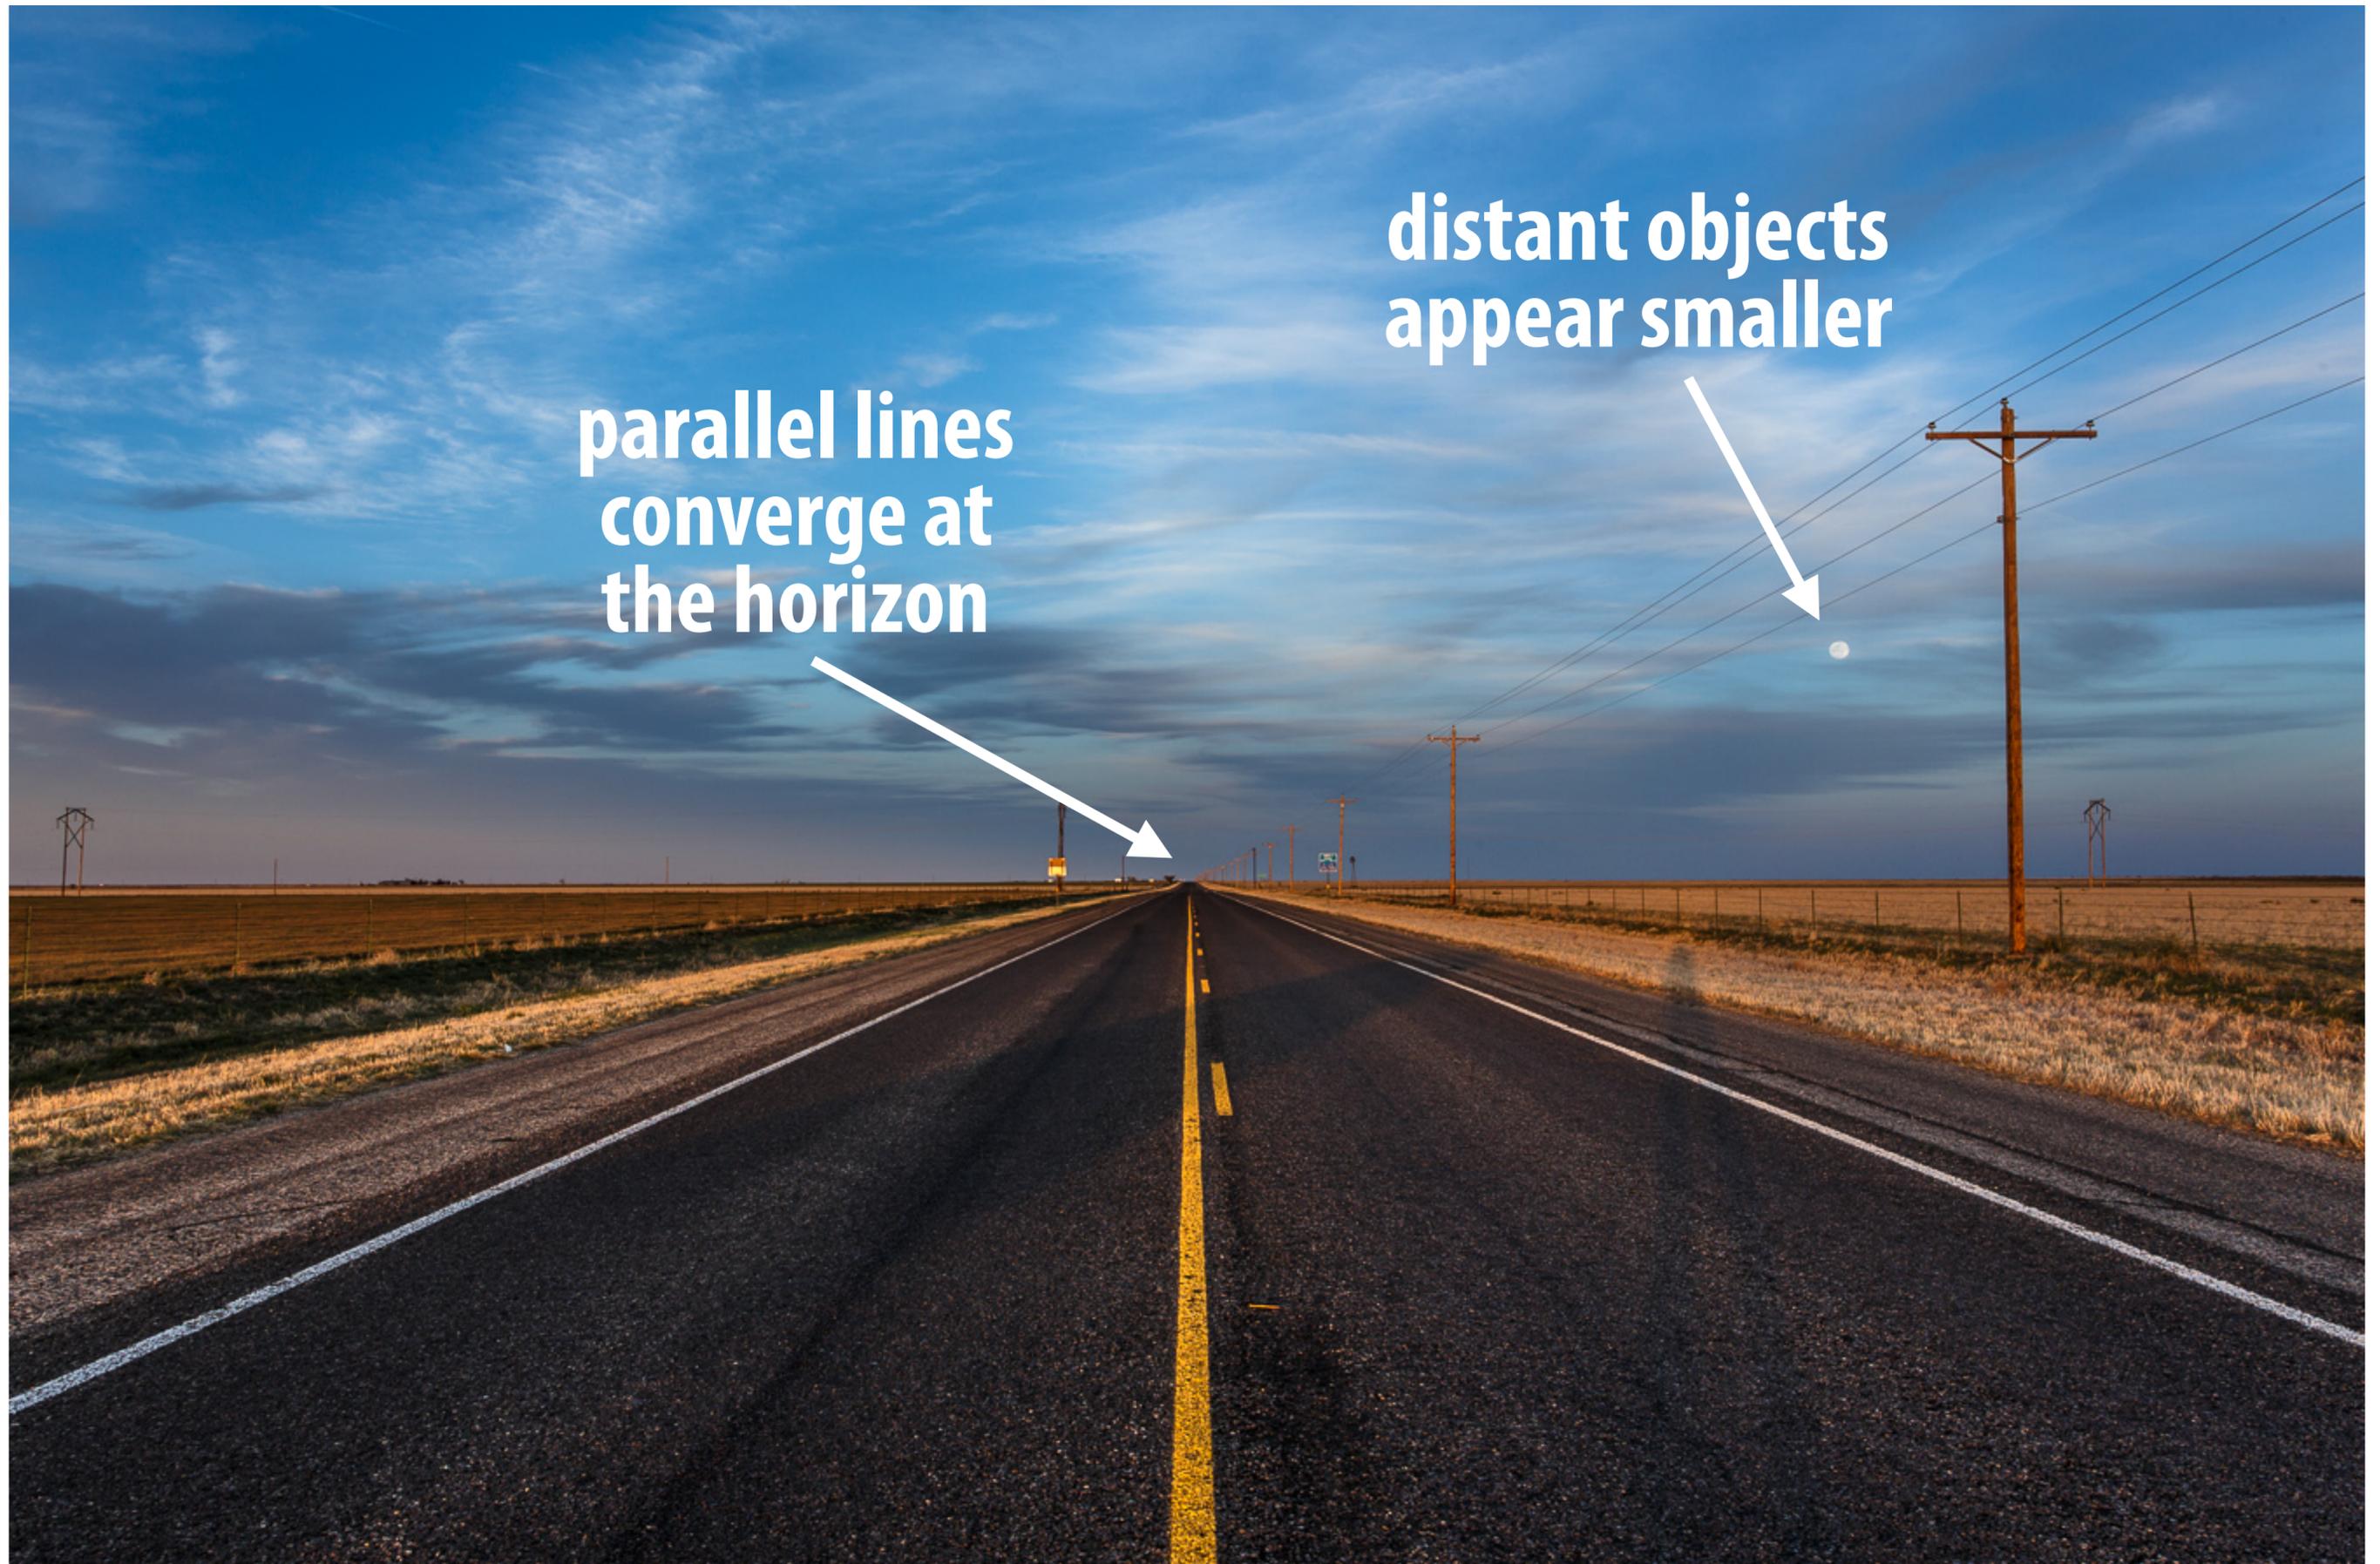
$$\mathbf{R}^T \mathbf{w} = [\mathbf{u} \cdot \mathbf{w} \quad \mathbf{v} \cdot \mathbf{w} \quad -\mathbf{w} \cdot \mathbf{w}]^T = [0 \quad 0 \quad -1]^T$$

Self-check exercise (for home)

- **Given a camera position P**
- **And a camera orientation given by orthonormal basis u, v, w (camera looking in w)**
- **What is a transformation matrix that places the scene in a coordinate space where...**
 - **The camera is at the origin**
 - **The camera is looking down $-z$.**

Perspective Projection

Perspective projection



Early painting: incorrect perspective



Carolingian painting from the 8-9th century

Perspective in art

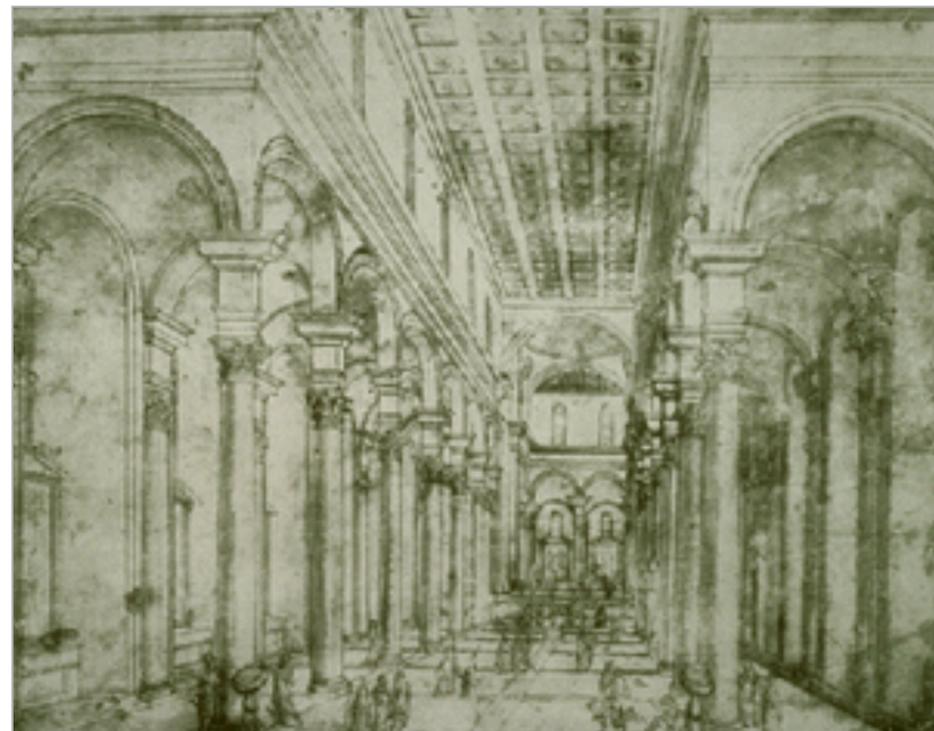


Giotto 1290

Evolution toward correct perspective



**Ambrogio Lorenzetti
Annunciation, 1344**



**Brunelleschi, elevation of Santo Spirito,
1434-83, Florence**



**Masaccio – The Tribute Money c.1426-27
Fresco, The Brancacci Chapel, Florence**

Perspective in art

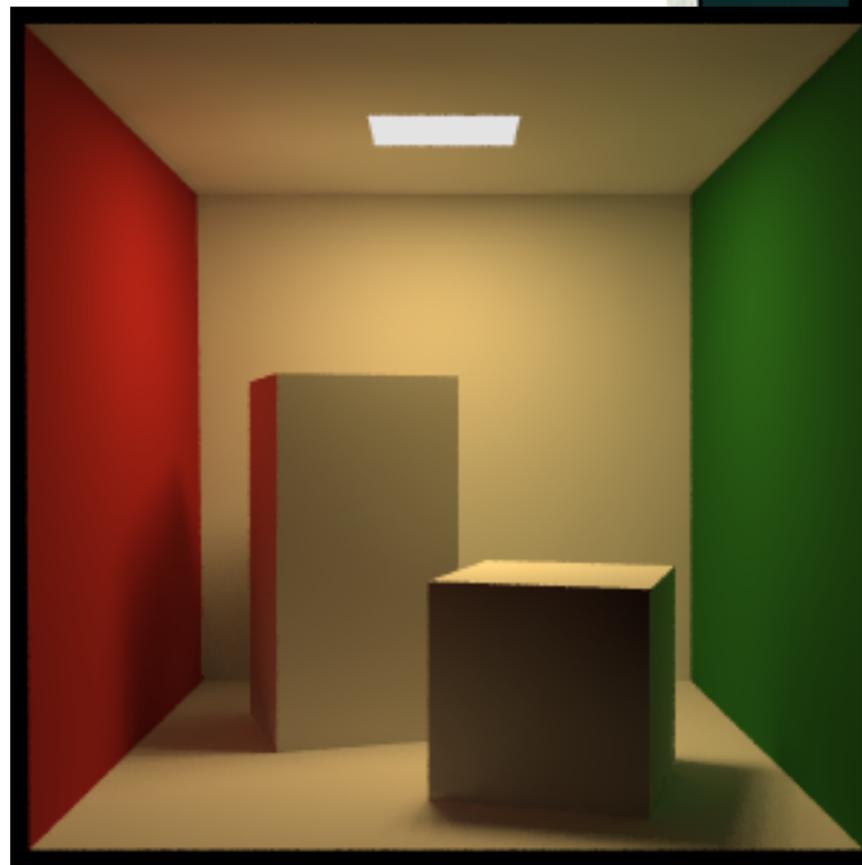
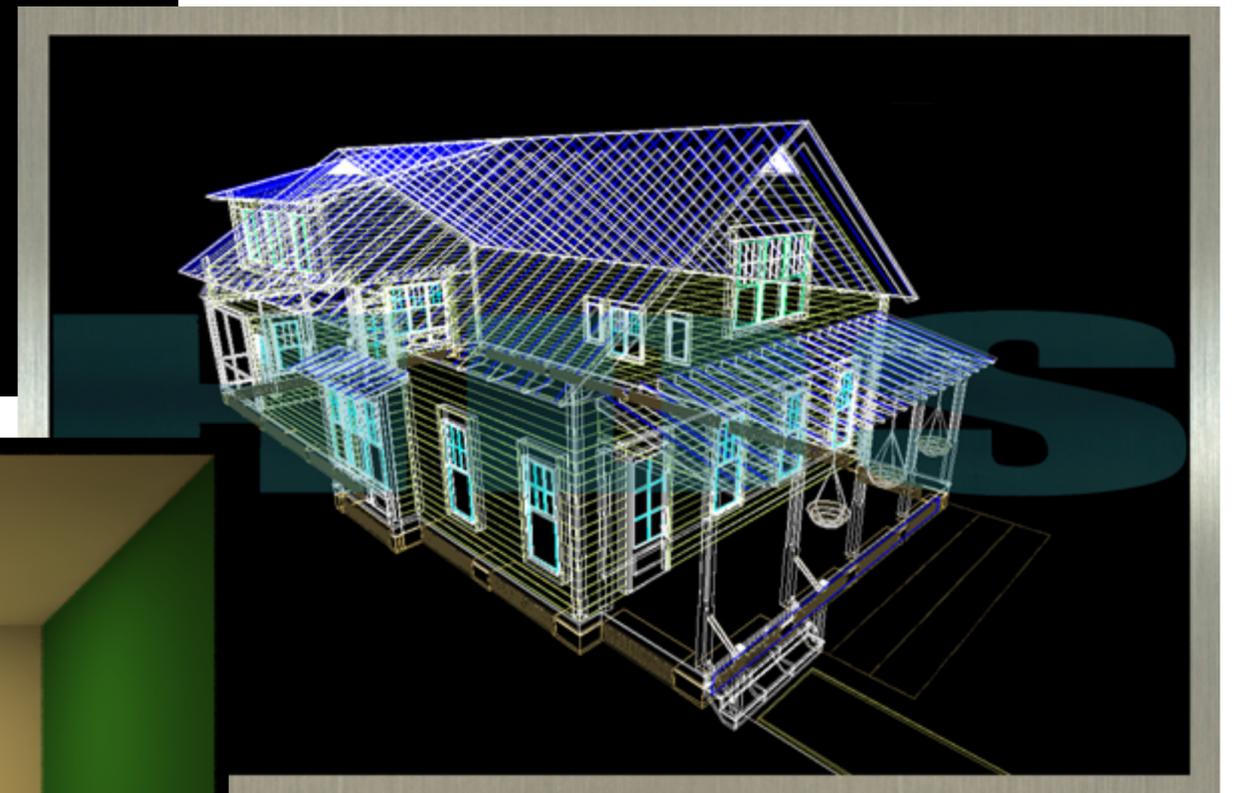
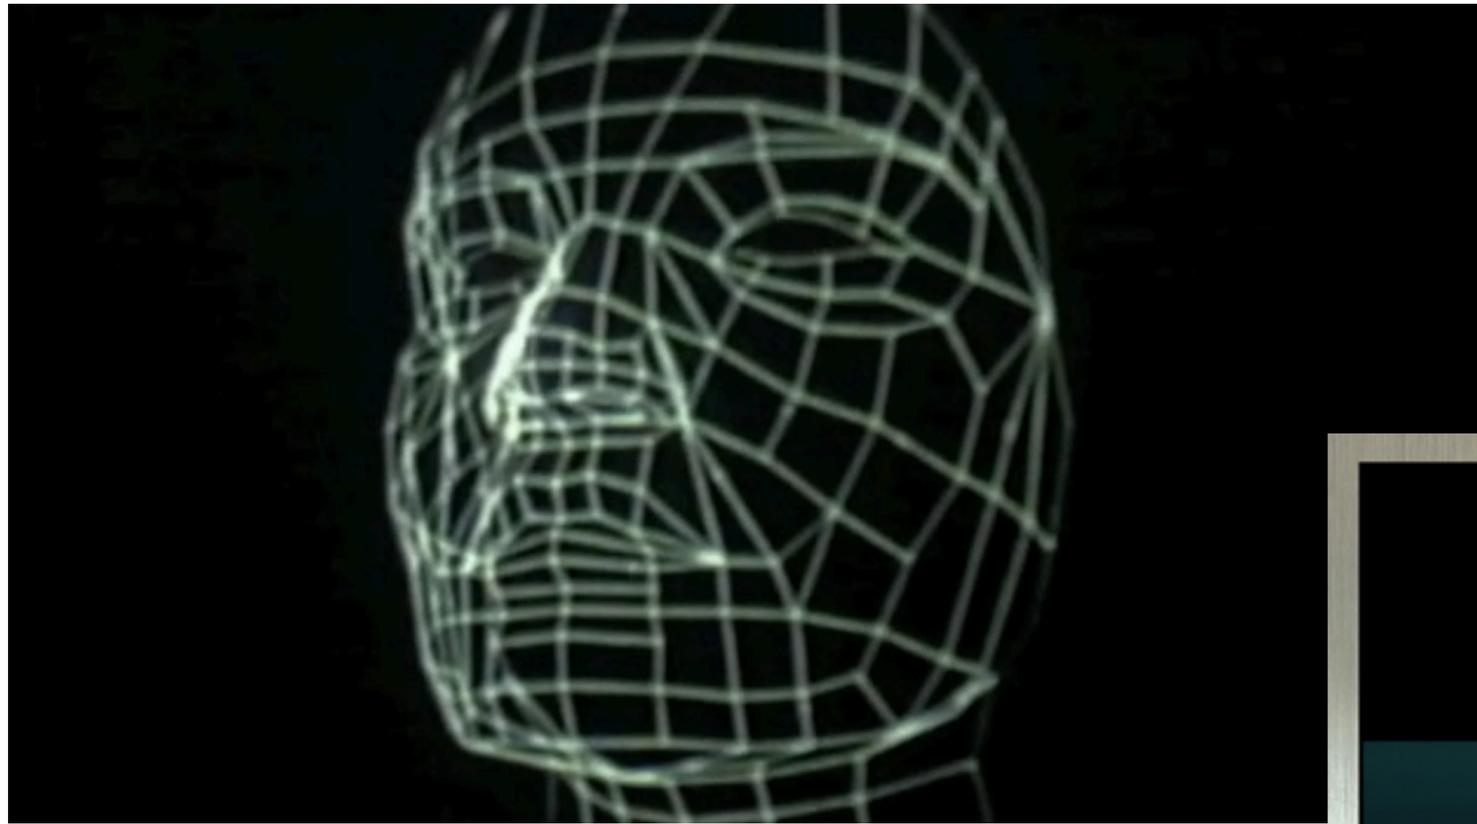


Delivery of the Keys (Sistine Chapel), Perugino, 1482

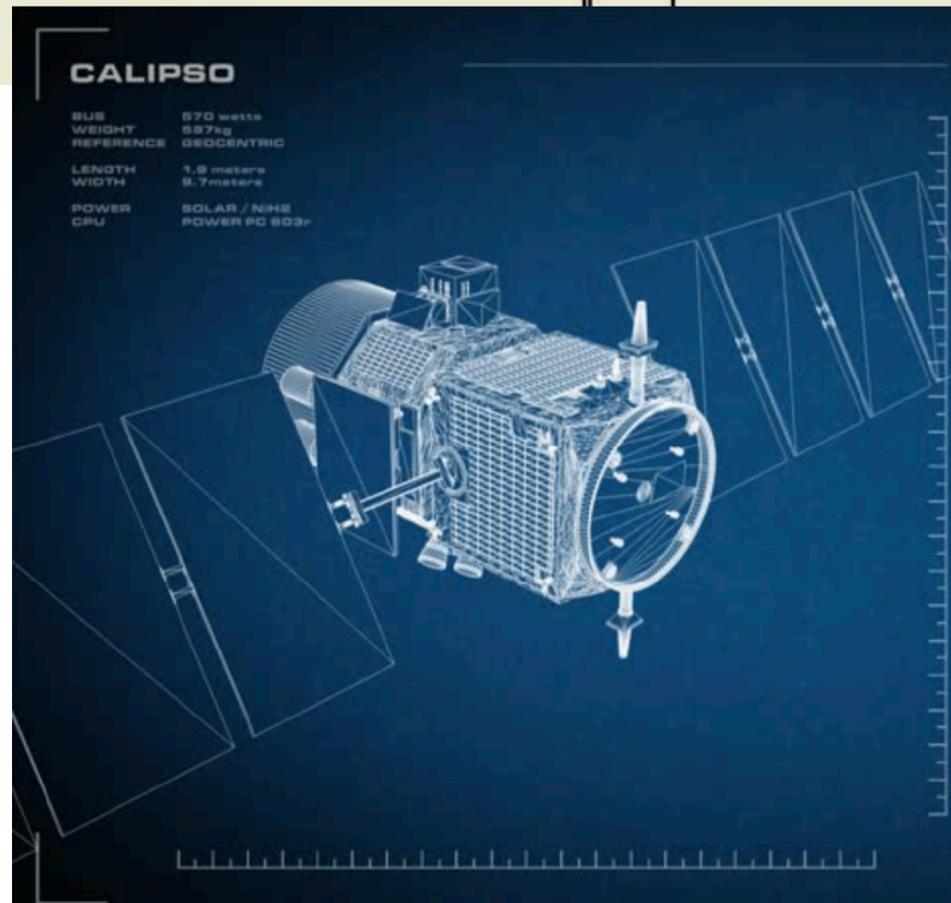
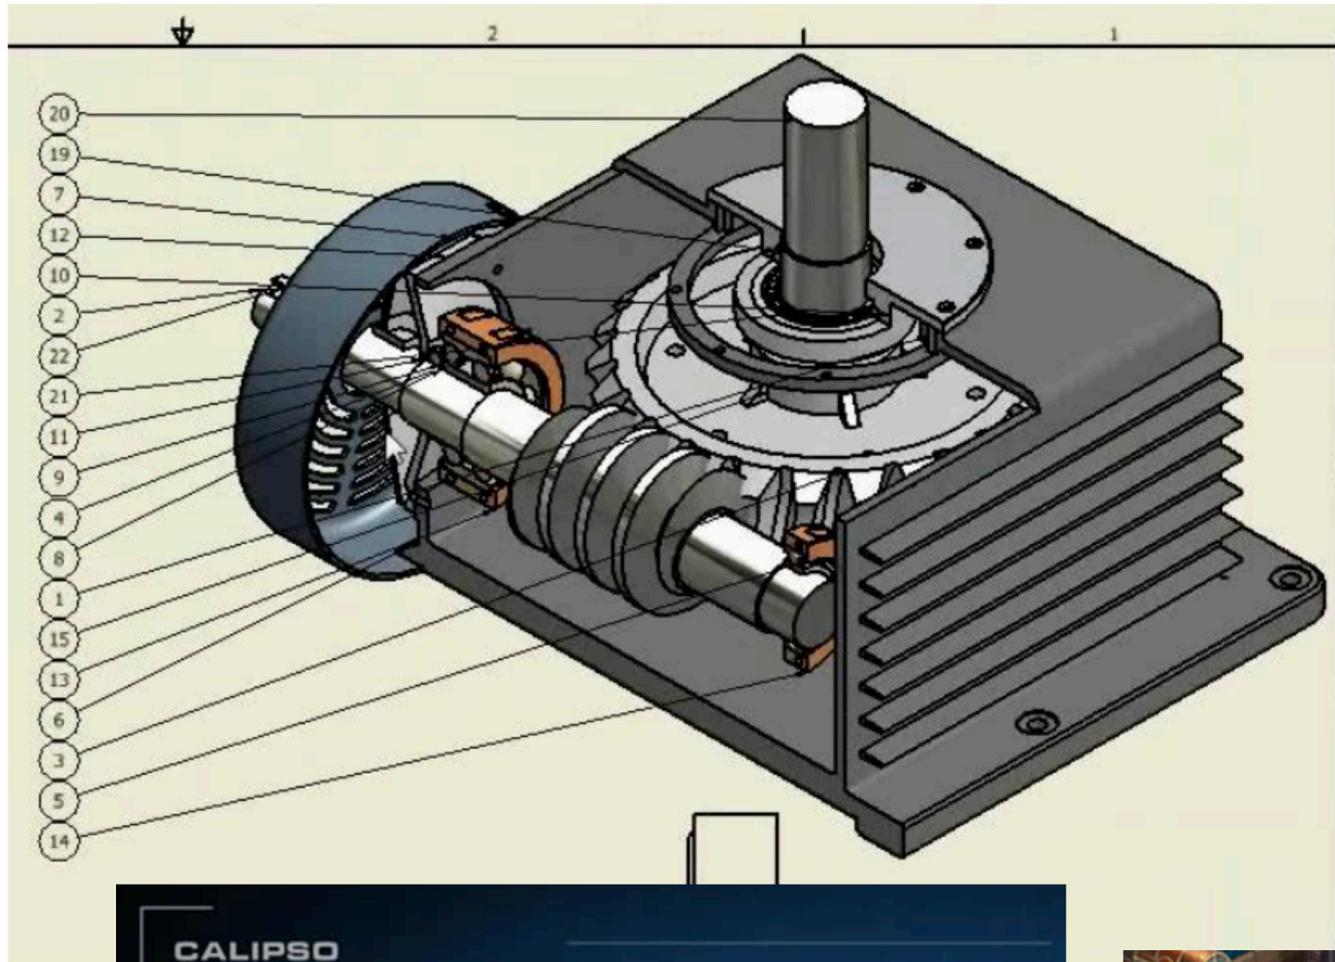
Later... rejection of proper perspective projection



Return of perspective in computer graphics

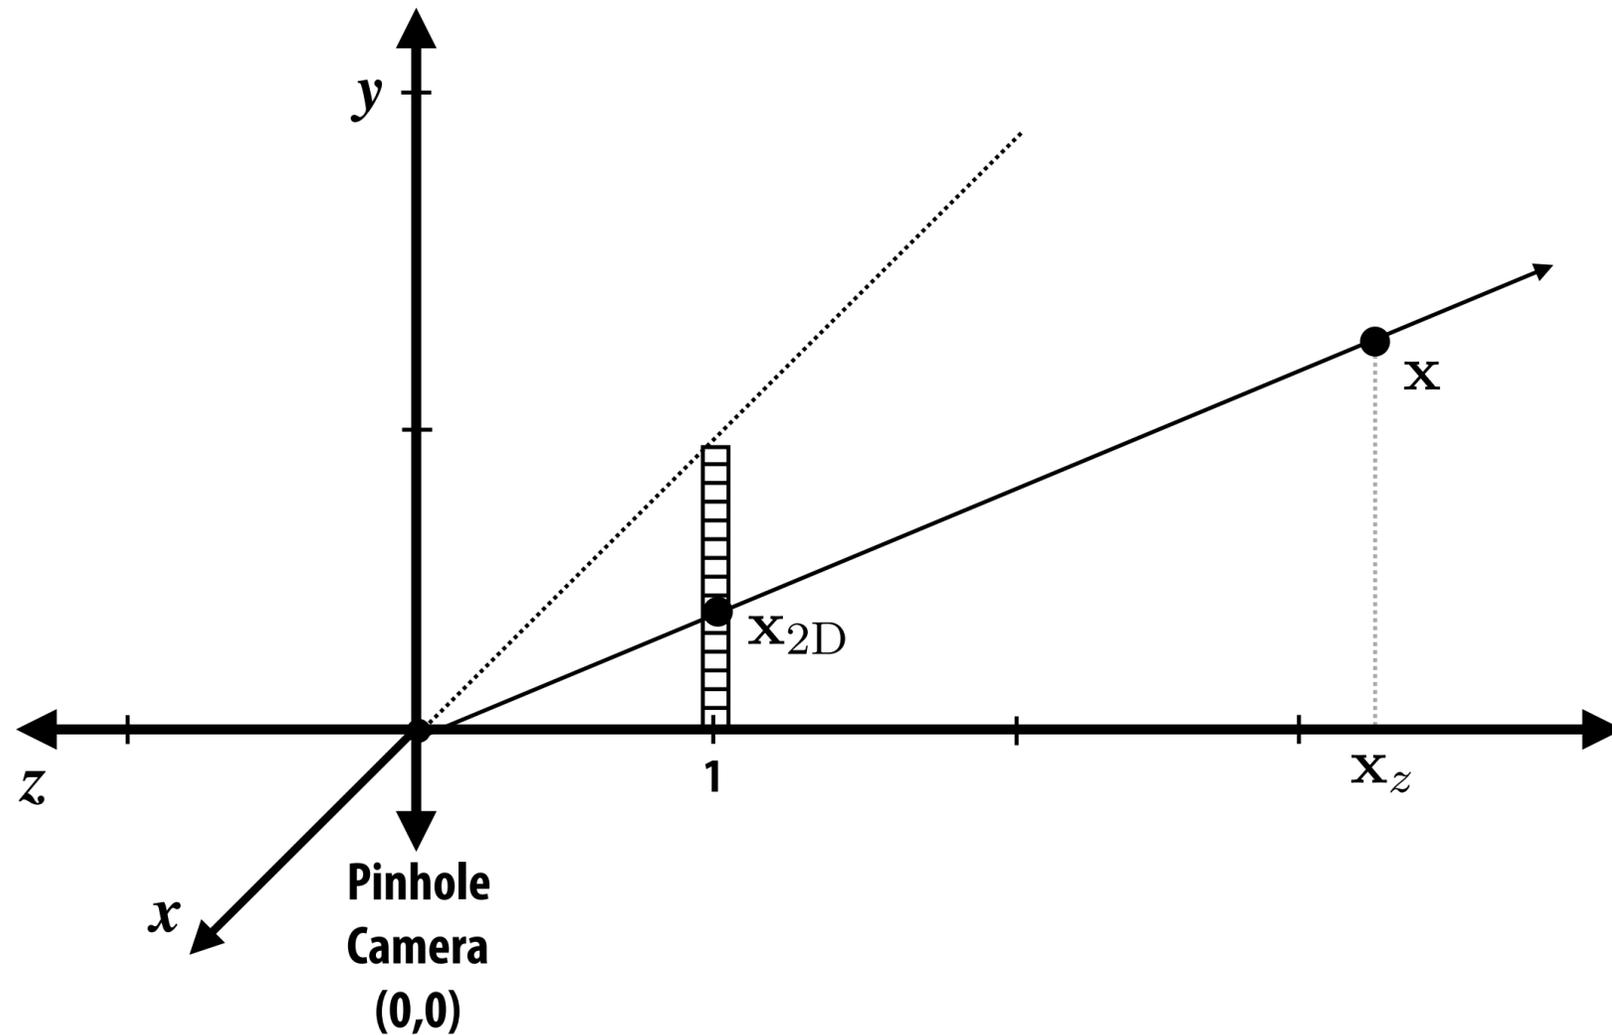


Rejection of perspective in computer graphics



Basic perspective projection

Input point in 3D-H: $\mathbf{x} = [x_x \quad x_y \quad x_z \quad 1]^T$



$$\mathbf{P} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

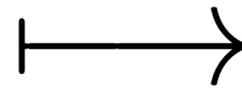
Assumption:

Pinhole camera at $(0,0)$ looking down $-z$

Perspective vs. orthographic projection

■ Most basic version of perspective matrix:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ w \end{bmatrix} = \begin{bmatrix} x \\ y \\ z \\ z \end{bmatrix}$$



$$\begin{bmatrix} x/z \\ y/z \\ 1 \\ 1 \end{bmatrix}$$

**objects shrink
in distance**

■ Most basic version of orthographic matrix:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$



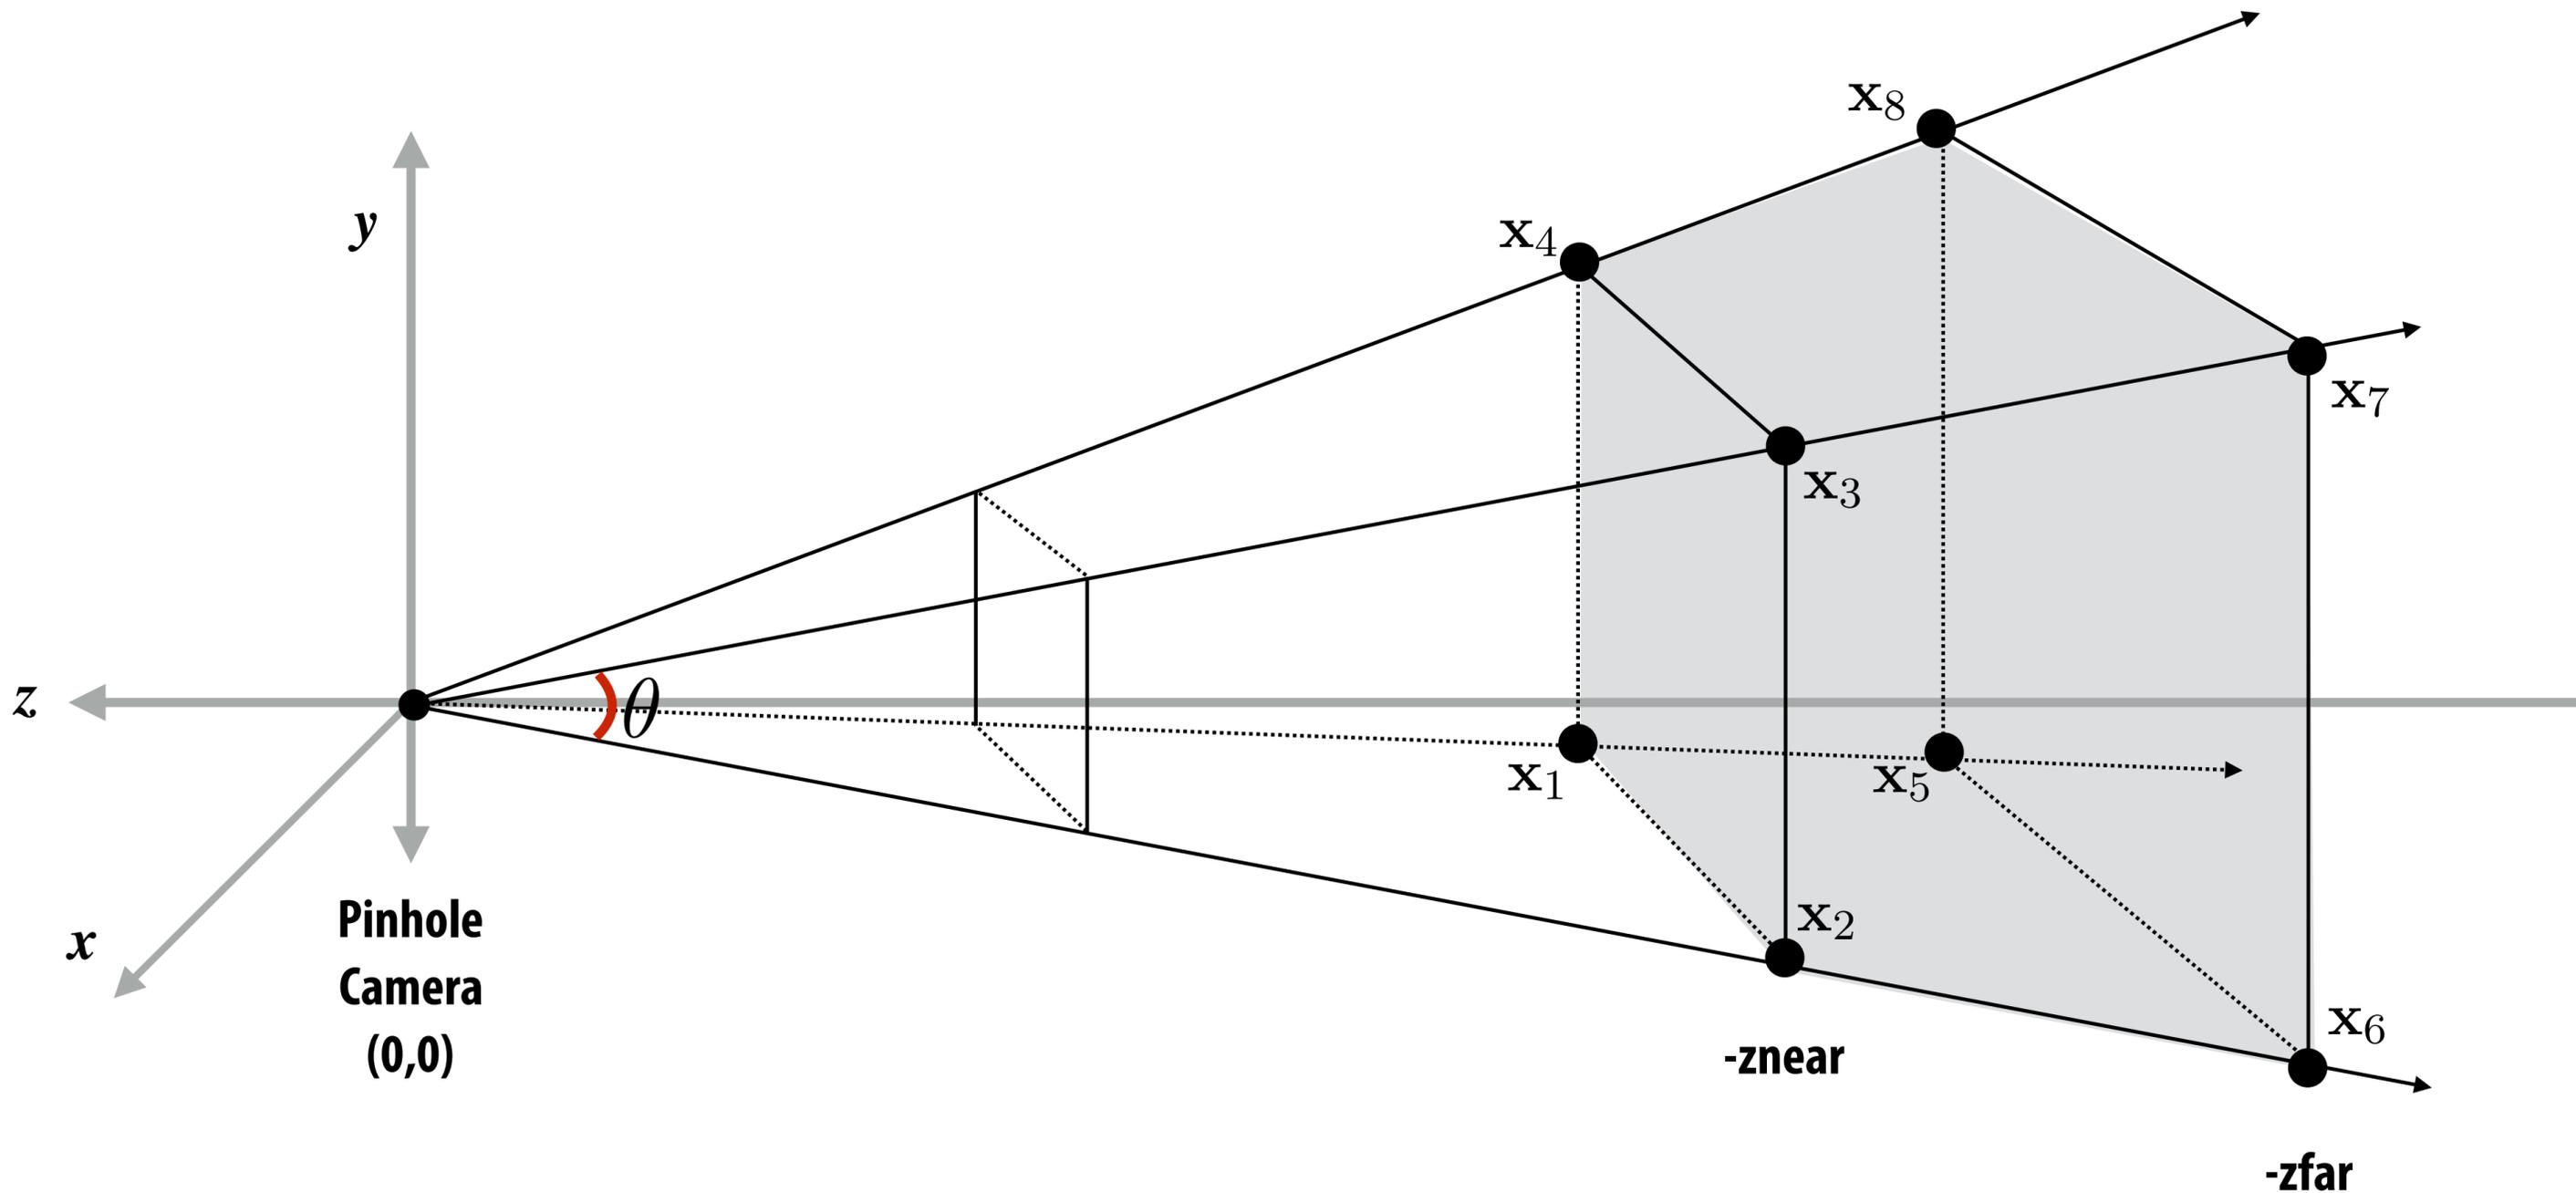
$$\begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

**objects stay the
same size**

...real projection matrices are a bit more complicated! :-)

View frustum

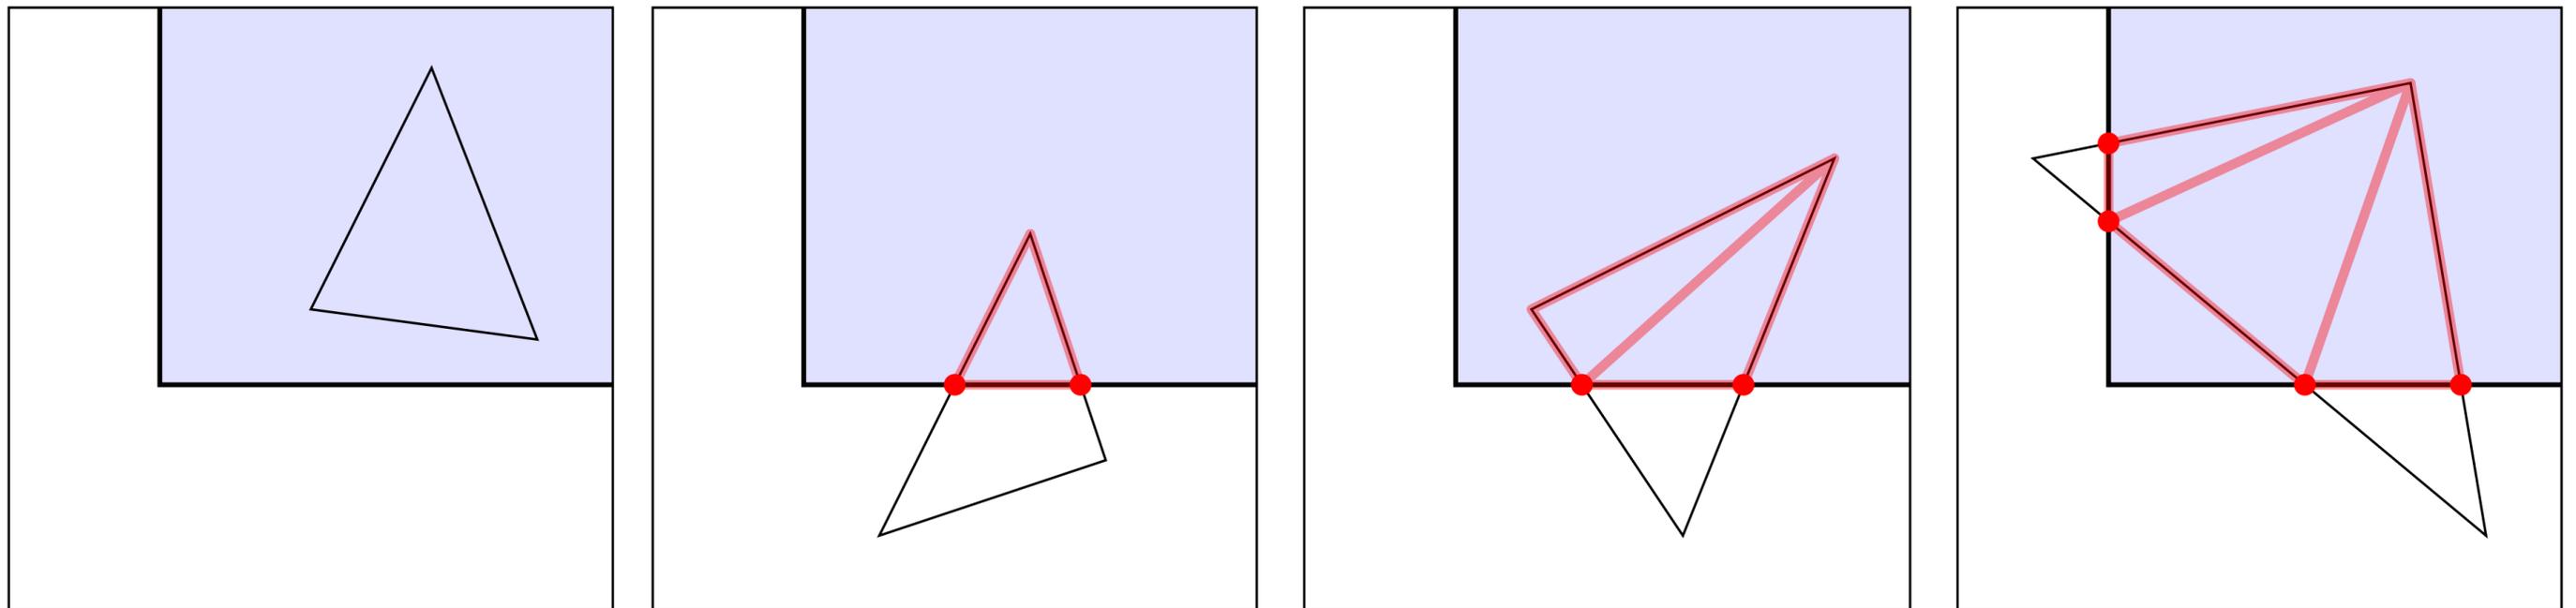
View frustum is the region of space the camera can see:



- Top/bottom/left/right planes correspond to sides of screen
- Near/far planes correspond to closest/furthest thing we want to draw

Clipping

- **“Clipping” is the process of eliminating triangles that aren’t visible to the camera (outside the view frustum)**
 - **Don’t waste time computing appearance of primitives you can’t see!**
 - **Sample-covered-by-triangle tests are expensive (“fine granularity” visibility)**
 - **Makes more sense to toss out entire primitives (“coarse granularity”)**
 - **Must deal with primitives that are partially clipped...**

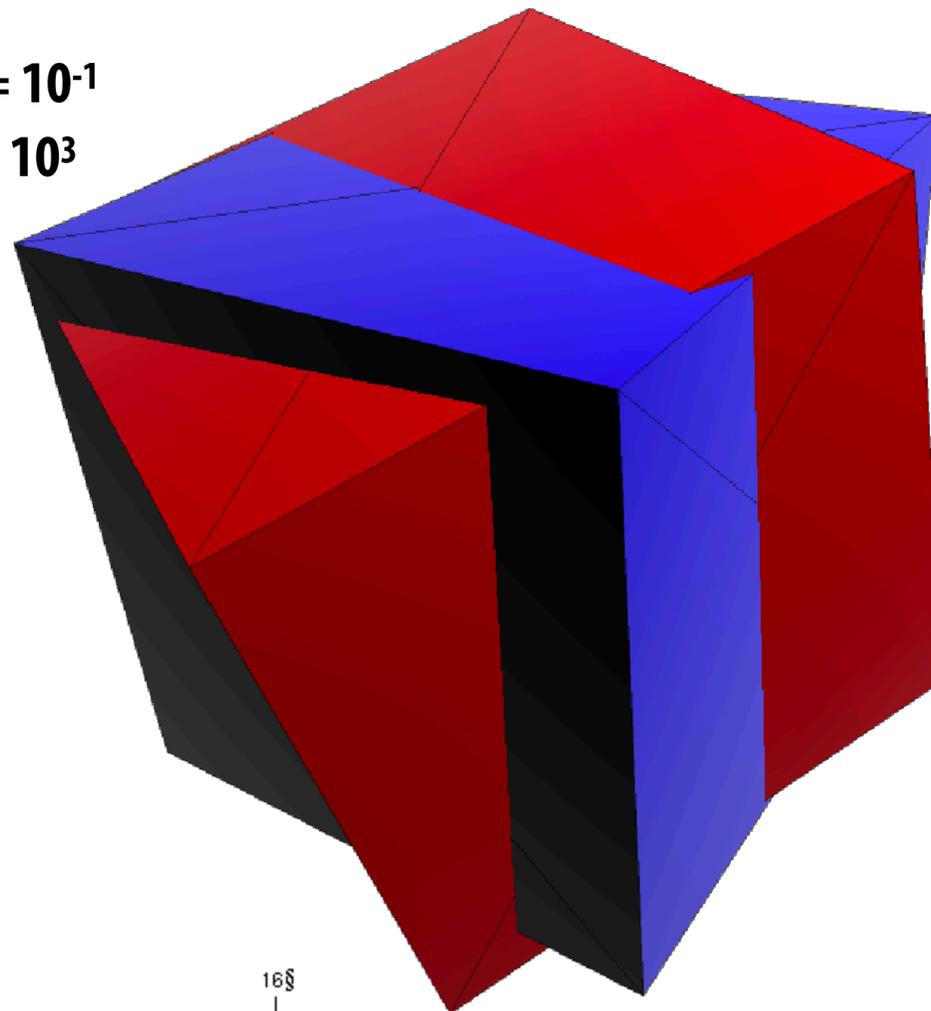


Aside: near/far plane clipping

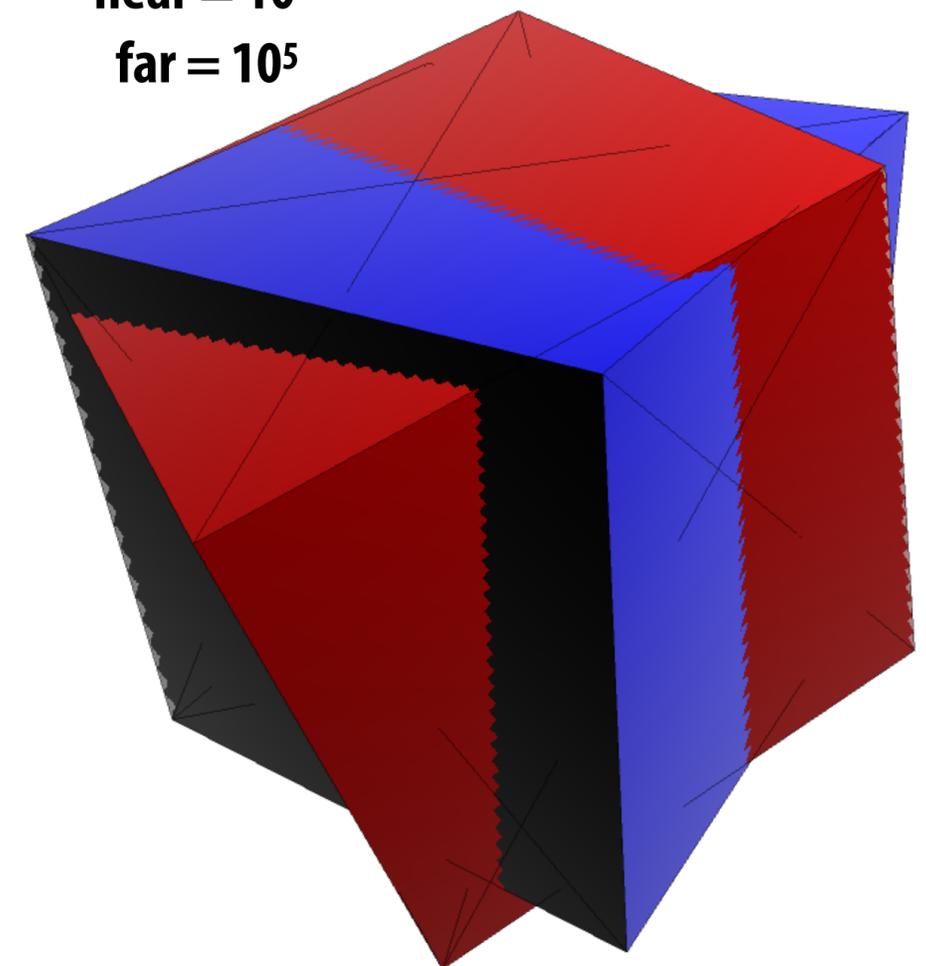
■ But why *near/far plane* clipping?

- Primitives (e.g., triangles) may have vertices both in front and behind camera!
(Causes headaches for rasterization, e.g., checking if fragments are behind camera)
- Also important for dealing with finite precision of depth buffer

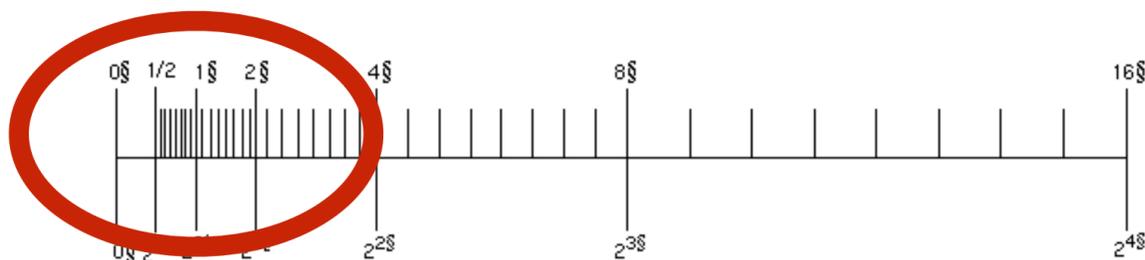
near = 10^{-1}
far = 10^3



near = 10^{-5}
far = 10^5



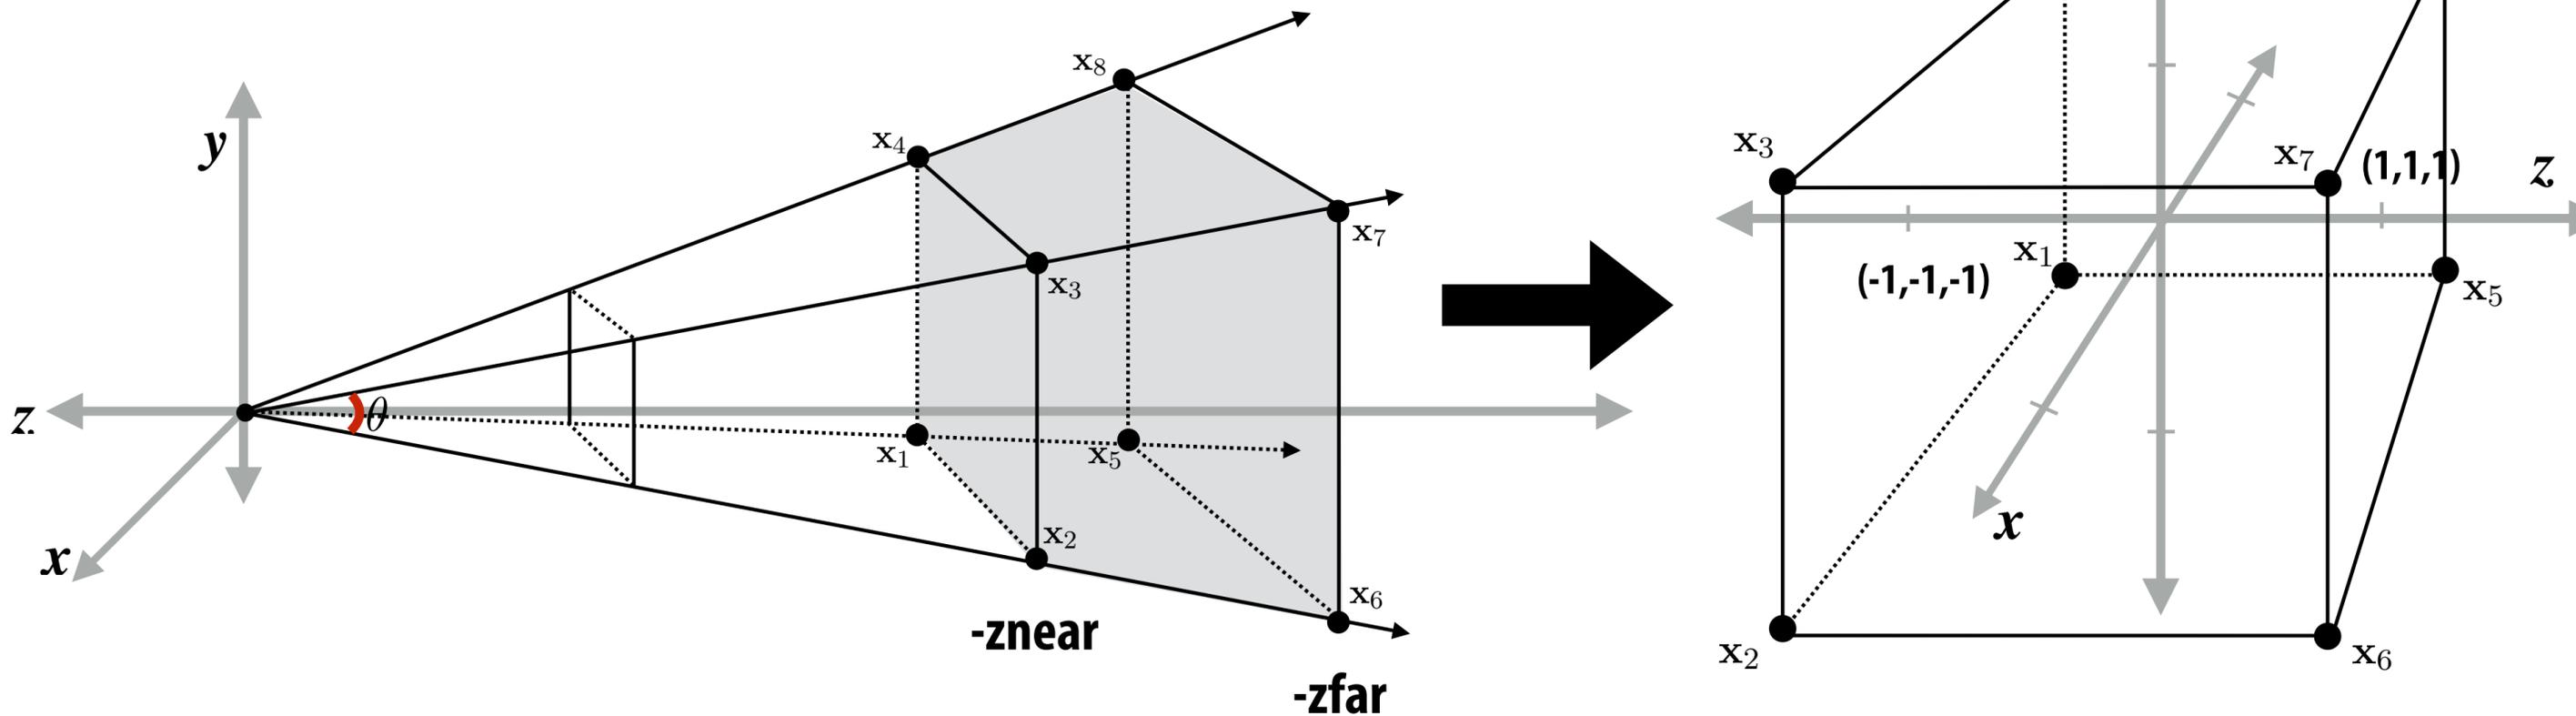
“Z-fighting”



floating point has more “resolution” near zero—hence more precise resolution of primitive-primitive intersection

Mapping frustum to normalized cube

Before mapping to 2D, map corners of frustum to corners of cube:



Why do we do this?

1. Makes *clipping* much easier!
 - can quickly discard geometry outside range $[-1, 1]$
2. Avoid issues of precision of perspective divide near origin
3. Different maps to cube yield different effects
 - specifically: perspective or orthographic
 - perspective is affine transformation, implemented via homogeneous coordinates
 - for orthographic view, just use identity matrix!

Perspective:

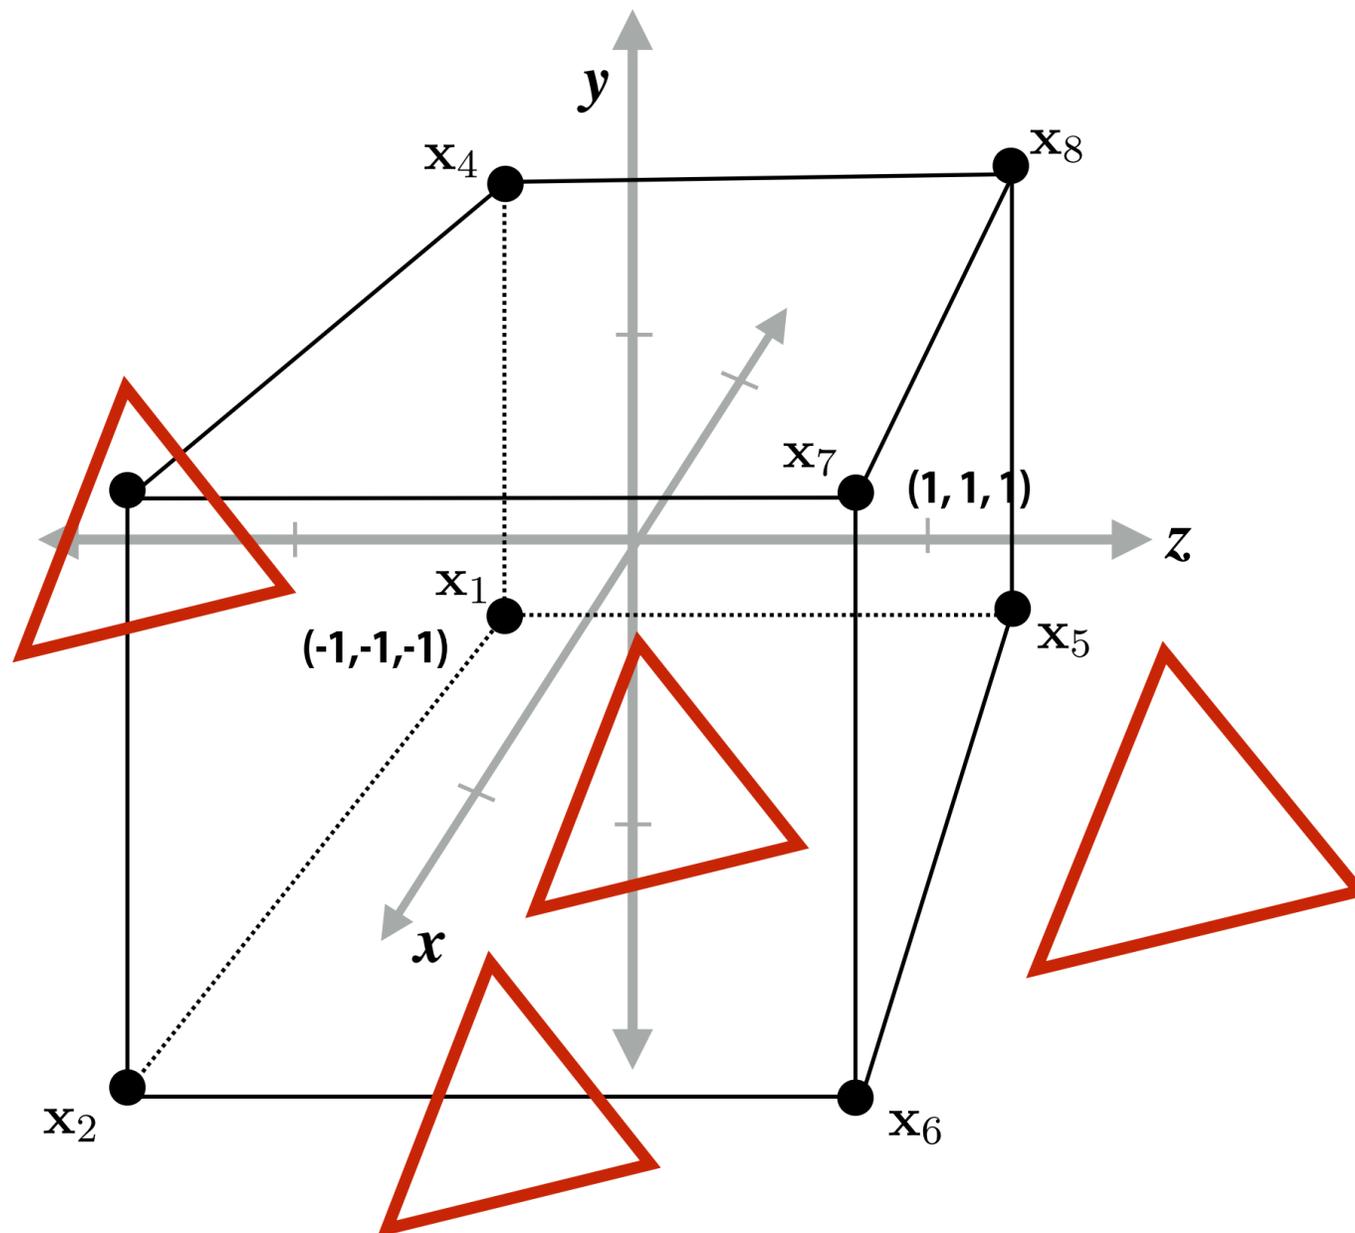
Set homogeneous coord to "z"
Distant objects get smaller

Orthographic: (not shown)

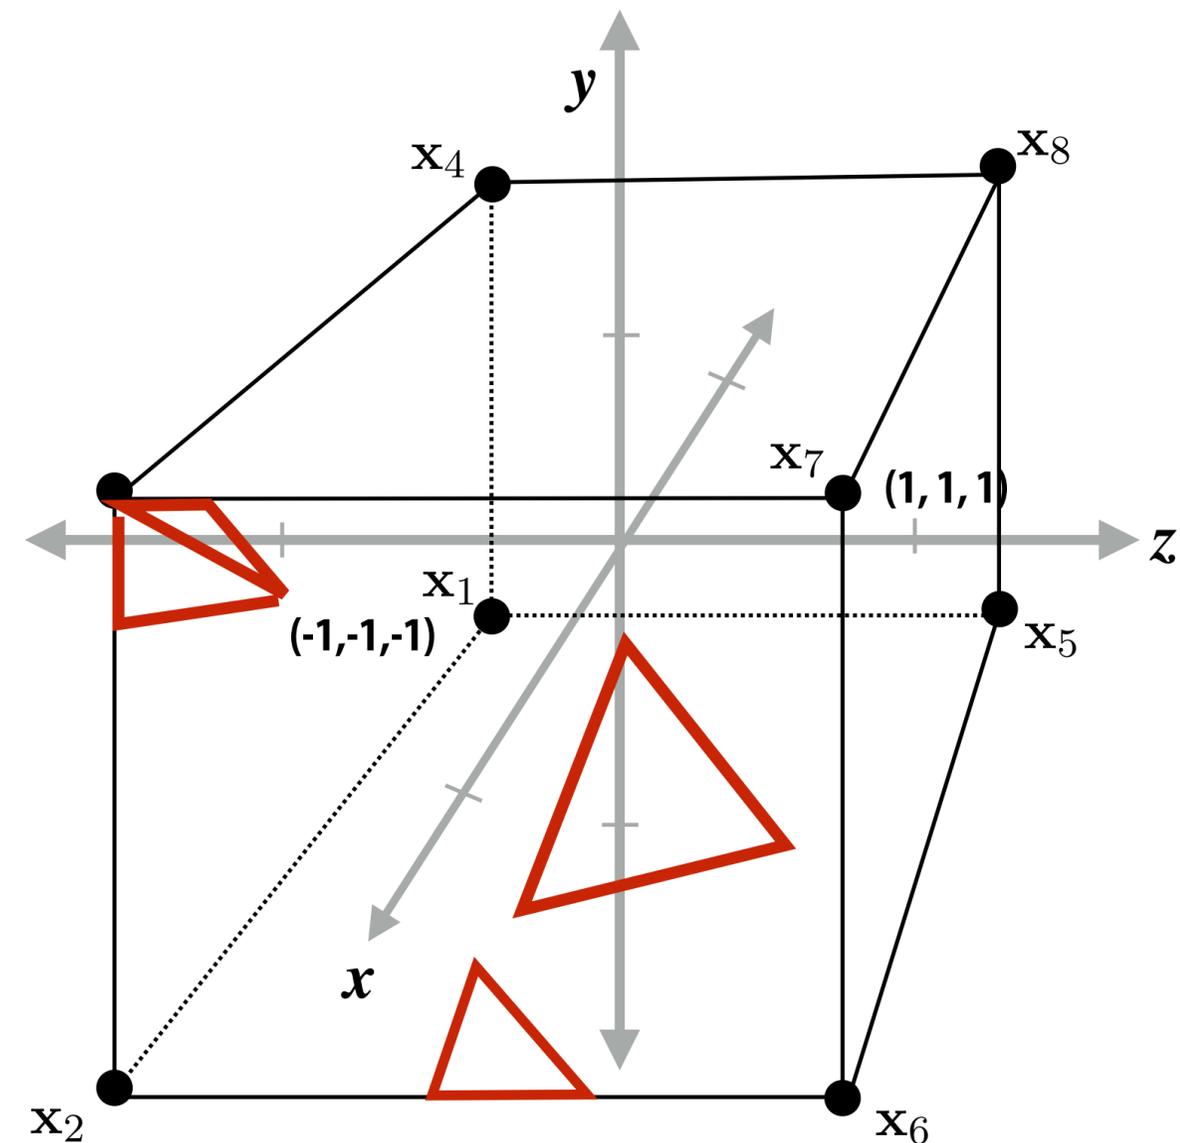
Set homogeneous coord to "1"
Distant objects remain same size

Clipping in normalized device coordinates (NDC)

- Discard triangles that lie complete outside the normalized cube (culling)
 - They are off screen, don't bother processing them further
- Clip triangles that extend beyond the cube... to the sides of the cube
 - Note: clipping may create more triangles



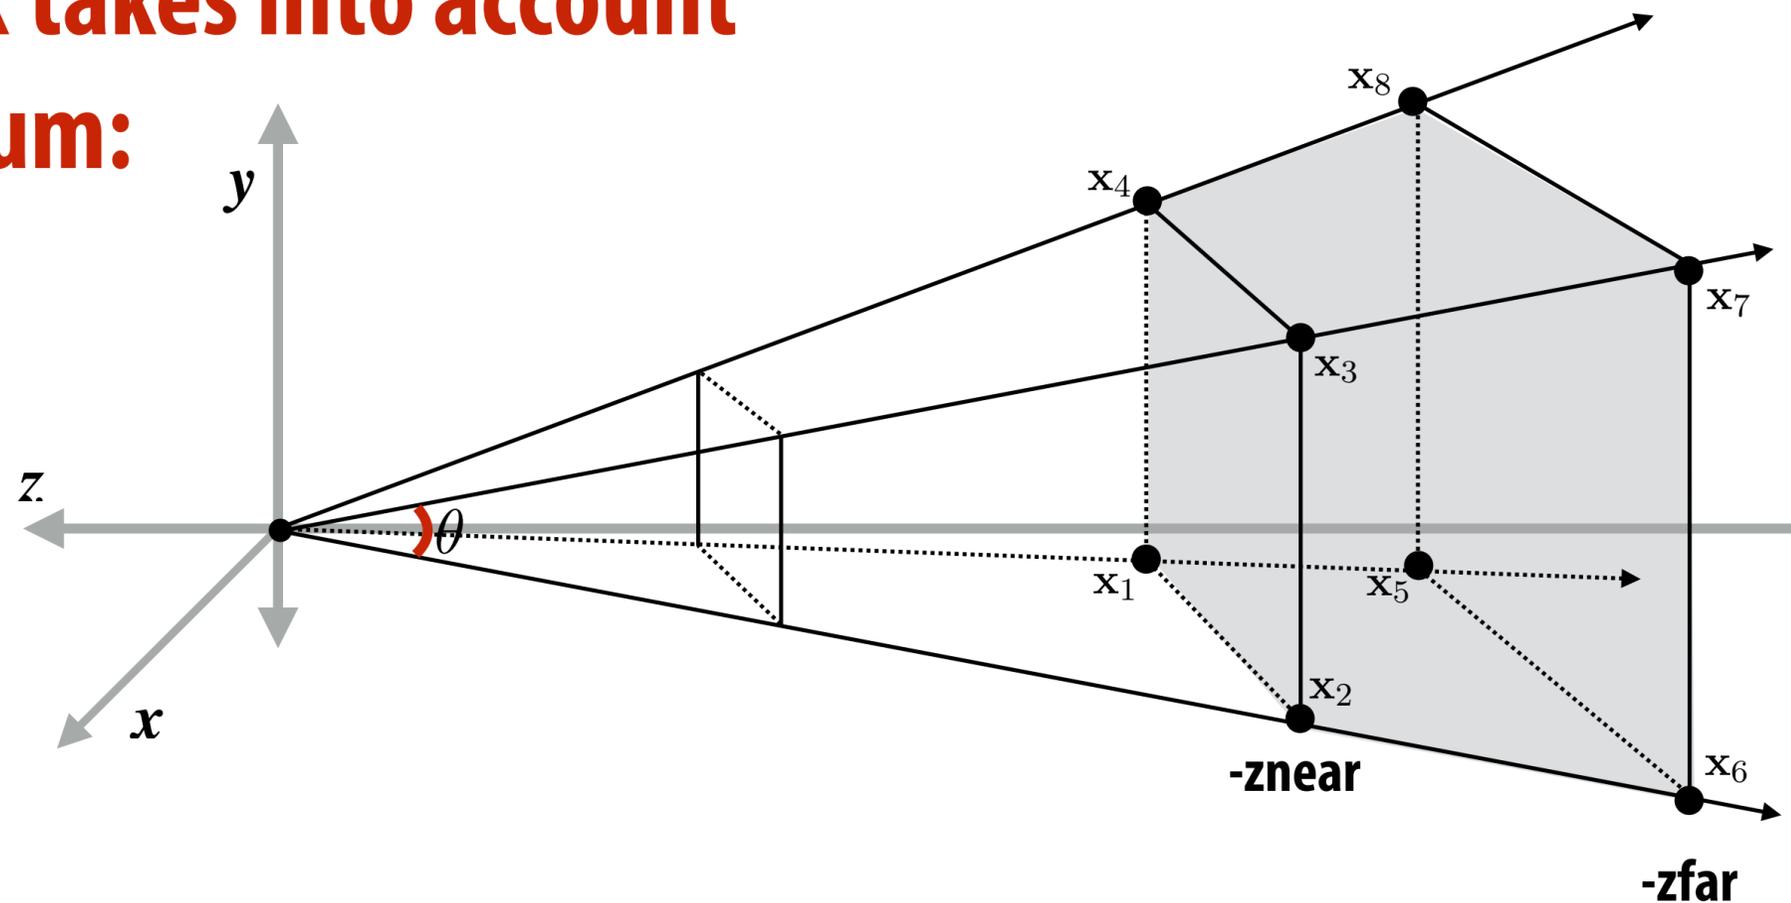
Triangles before clipping



Triangles after clipping

Matrix for perspective transform

Real perspective matrix takes into account geometry of view frustum:



$$\begin{pmatrix} \frac{n}{r} & 0 & 0 & 0 \\ 0 & \frac{n}{t} & 0 & 0 \\ 0 & 0 & \frac{-(f+n)}{f-n} & \frac{-2fn}{f-n} \\ 0 & 0 & -1 & 0 \end{pmatrix}$$

left (l), right (r), top (t), bottom (b), near (n), far (f)

(matrix at left is perspective projection for frustum that is symmetric about x,y axes: l=-r, t=-b)

Review: screen transform

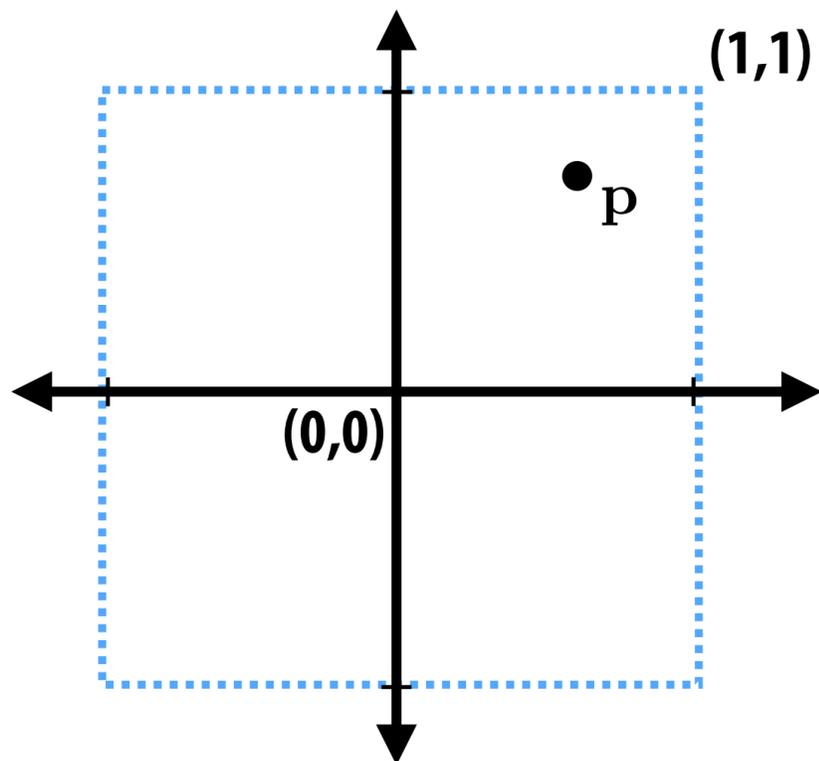
After divide, coordinates in $[-1,1]$ have to be “stretched” to fit the screen

Example:

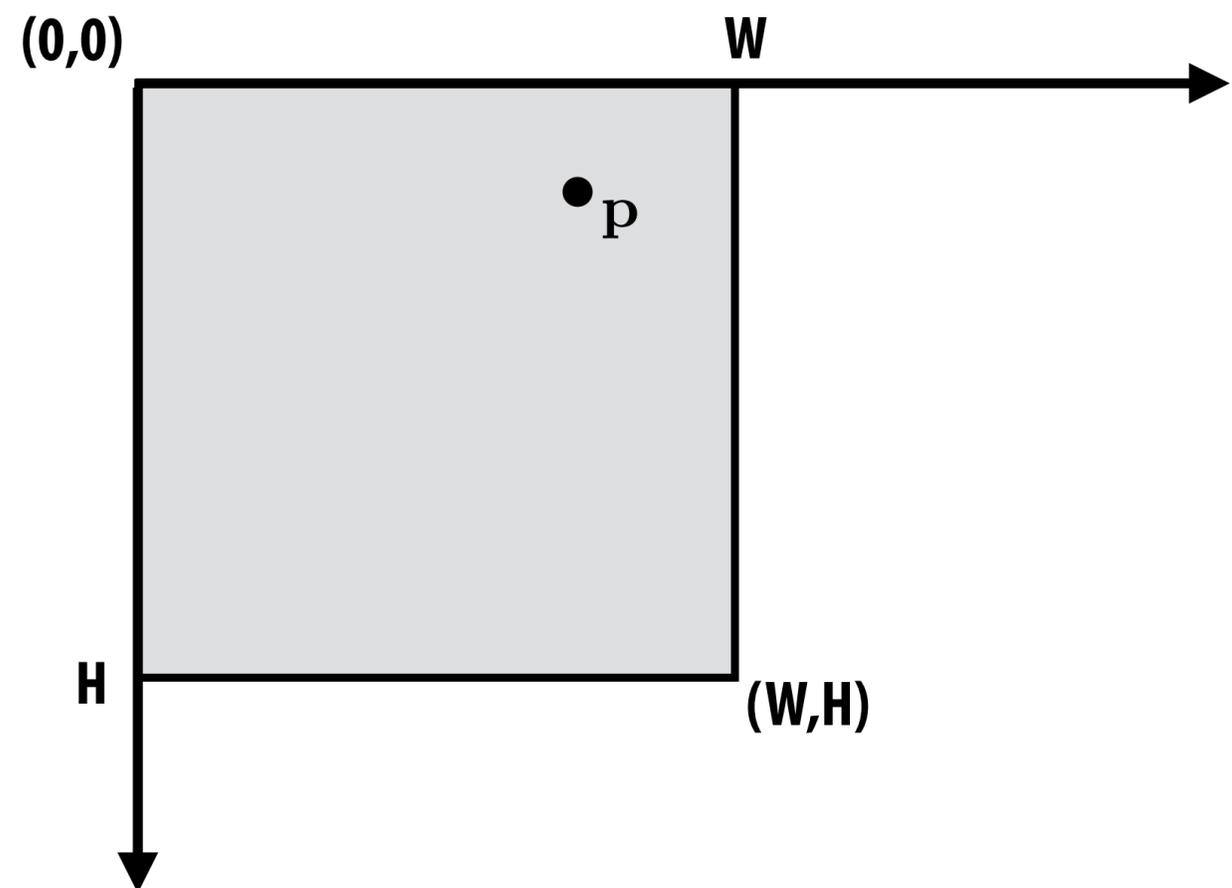
All points within $(-1,1)$ to $(1,1)$ region are on screen

$(1,1)$ in normalized space maps to $(W,0)$ in screen

Normalized coordinate space:



Screen ($W \times H$ output image) coordinate space:



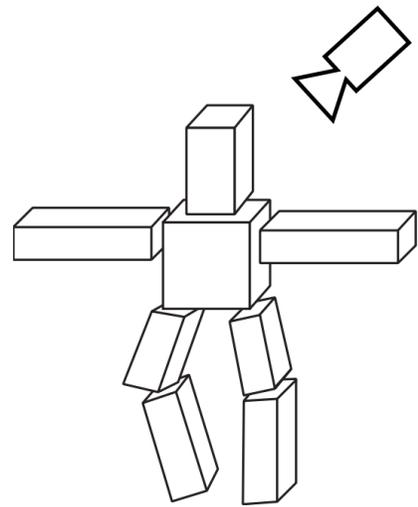
Step 1: reflect about x-axis

Step 2: translate by $(1,1)$

Step 3: scale by $(W/2, H/2)$

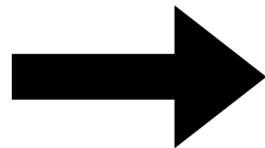
Transformations: from objects to the screen

[WORLD COORDINATES]

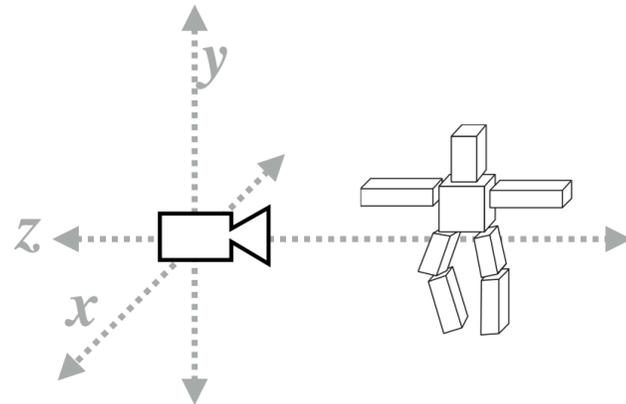


original description
of objects

view
transform

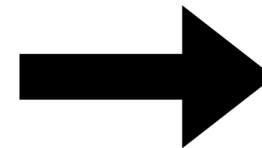


[VIEW COORDINATES]

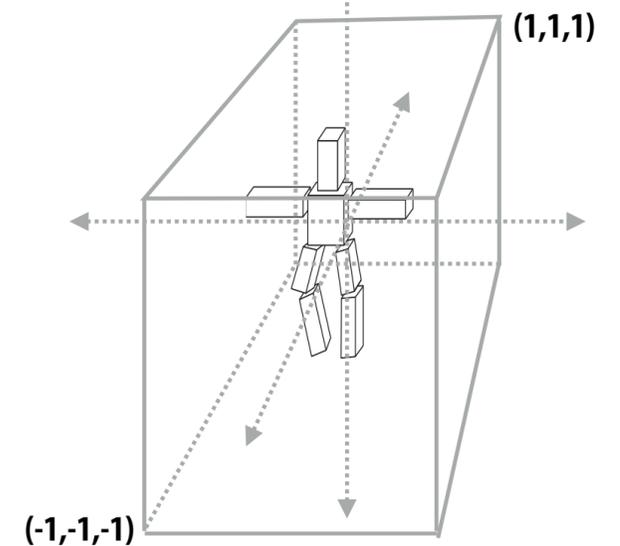


all positions now expressed
relative to camera; camera is
sitting at origin looking down
-z direction
(can canonicalize projection)

projection
transform

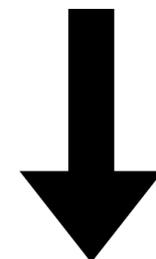


[CLIP COORDINATES]

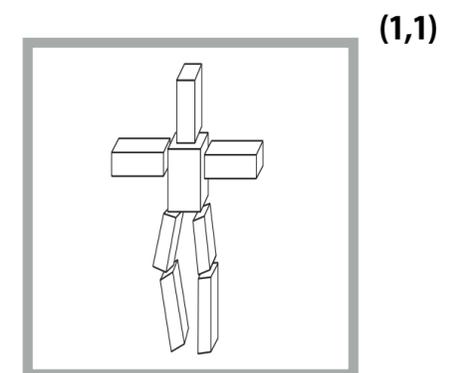


everything visible to the
camera is mapped to unit
cube for easy "clipping"

perspective
divide

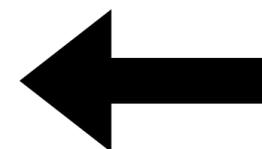


[NORMALIZED COORDINATES]

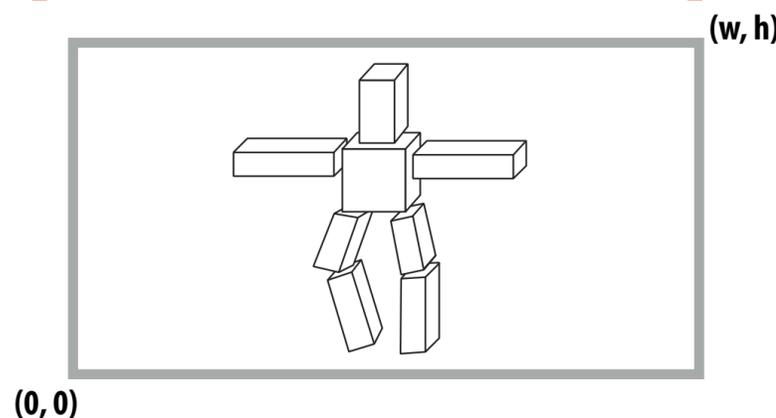


unit cube mapped to unit
square via perspective divide

screen
transform

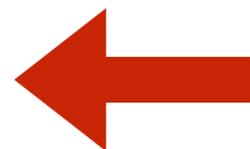


[WINDOW COORDINATES]



Screen transform:
objects now in 2D screen coordinates

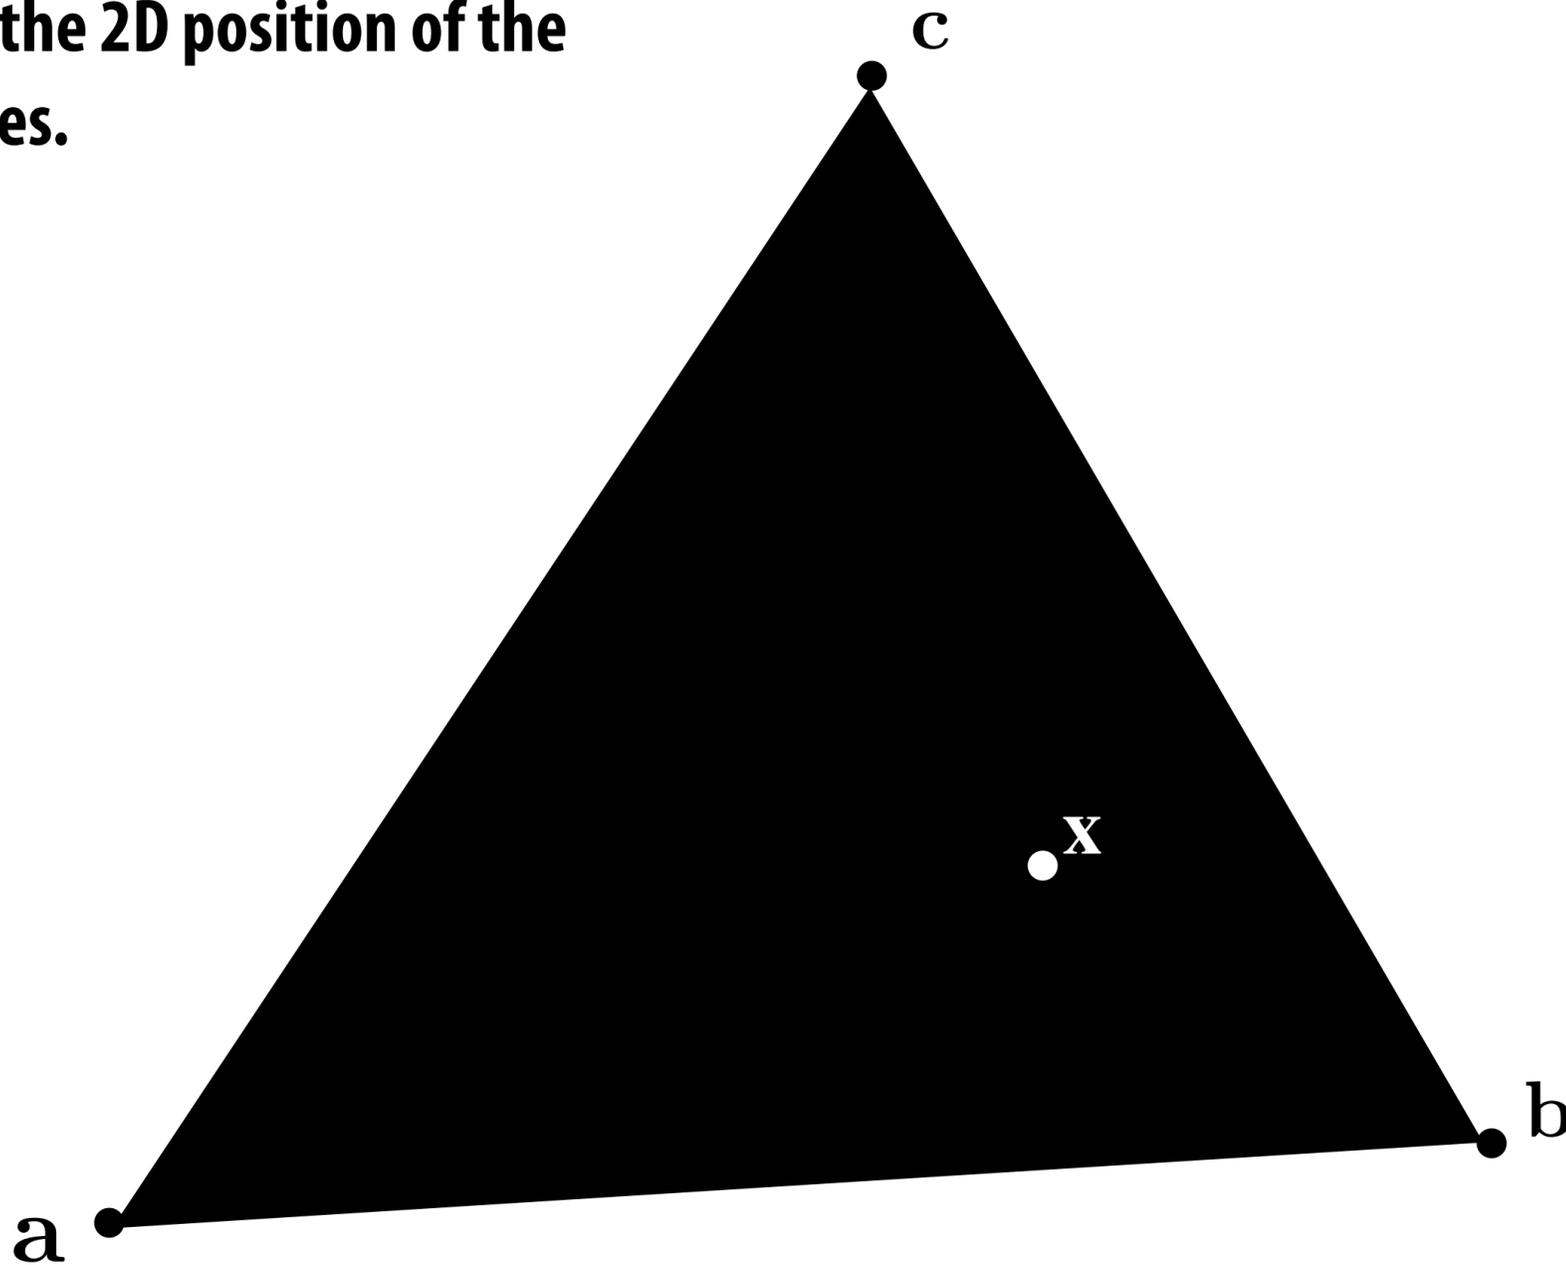
primitives are now 2D
and can be drawn via
rasterization



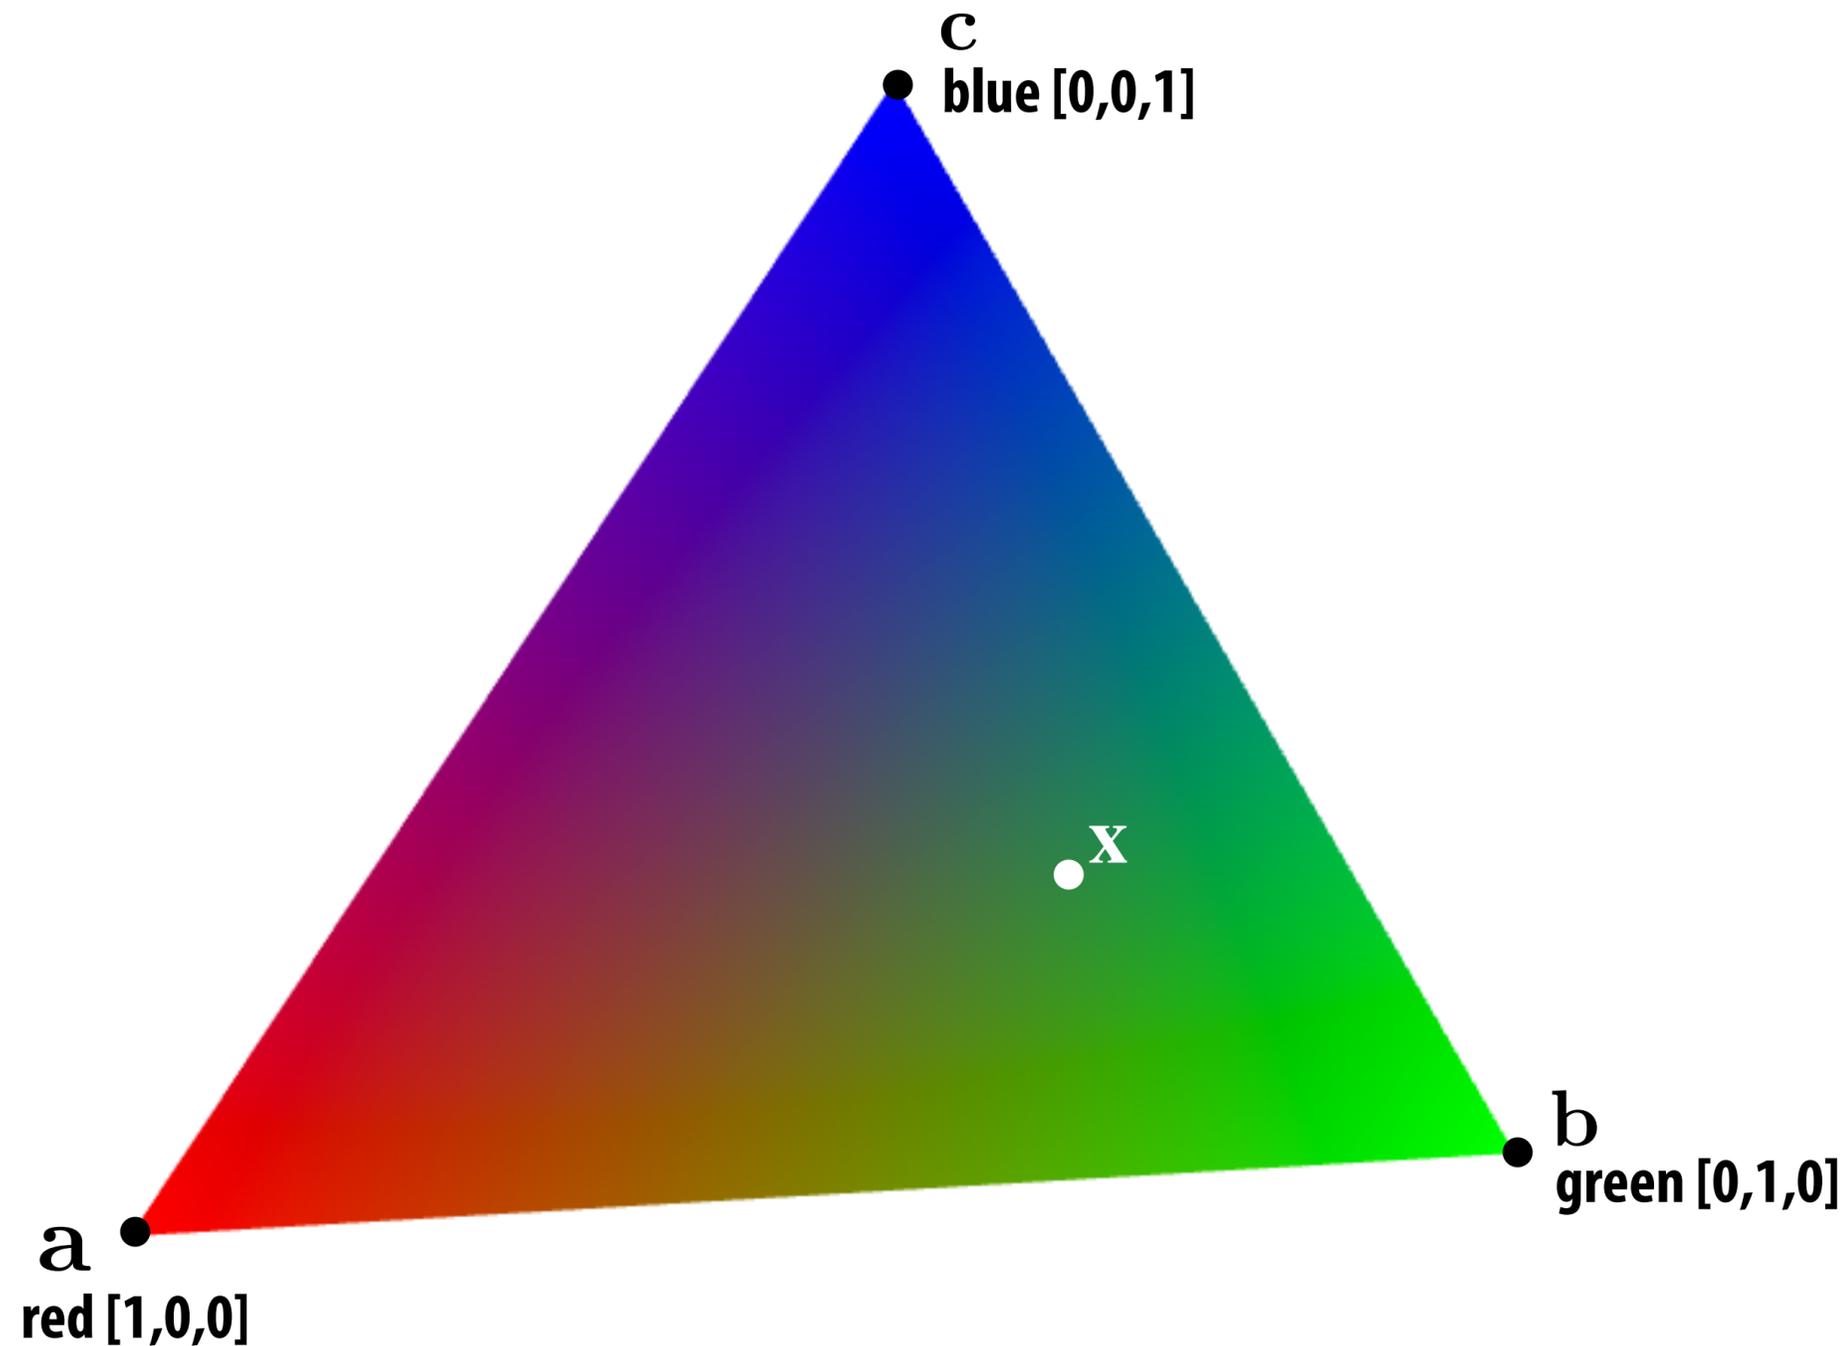
Texture mapping

Coverage(x, y)

In lecture 2 we discussed how to sample coverage given the 2D position of the triangle's vertices.



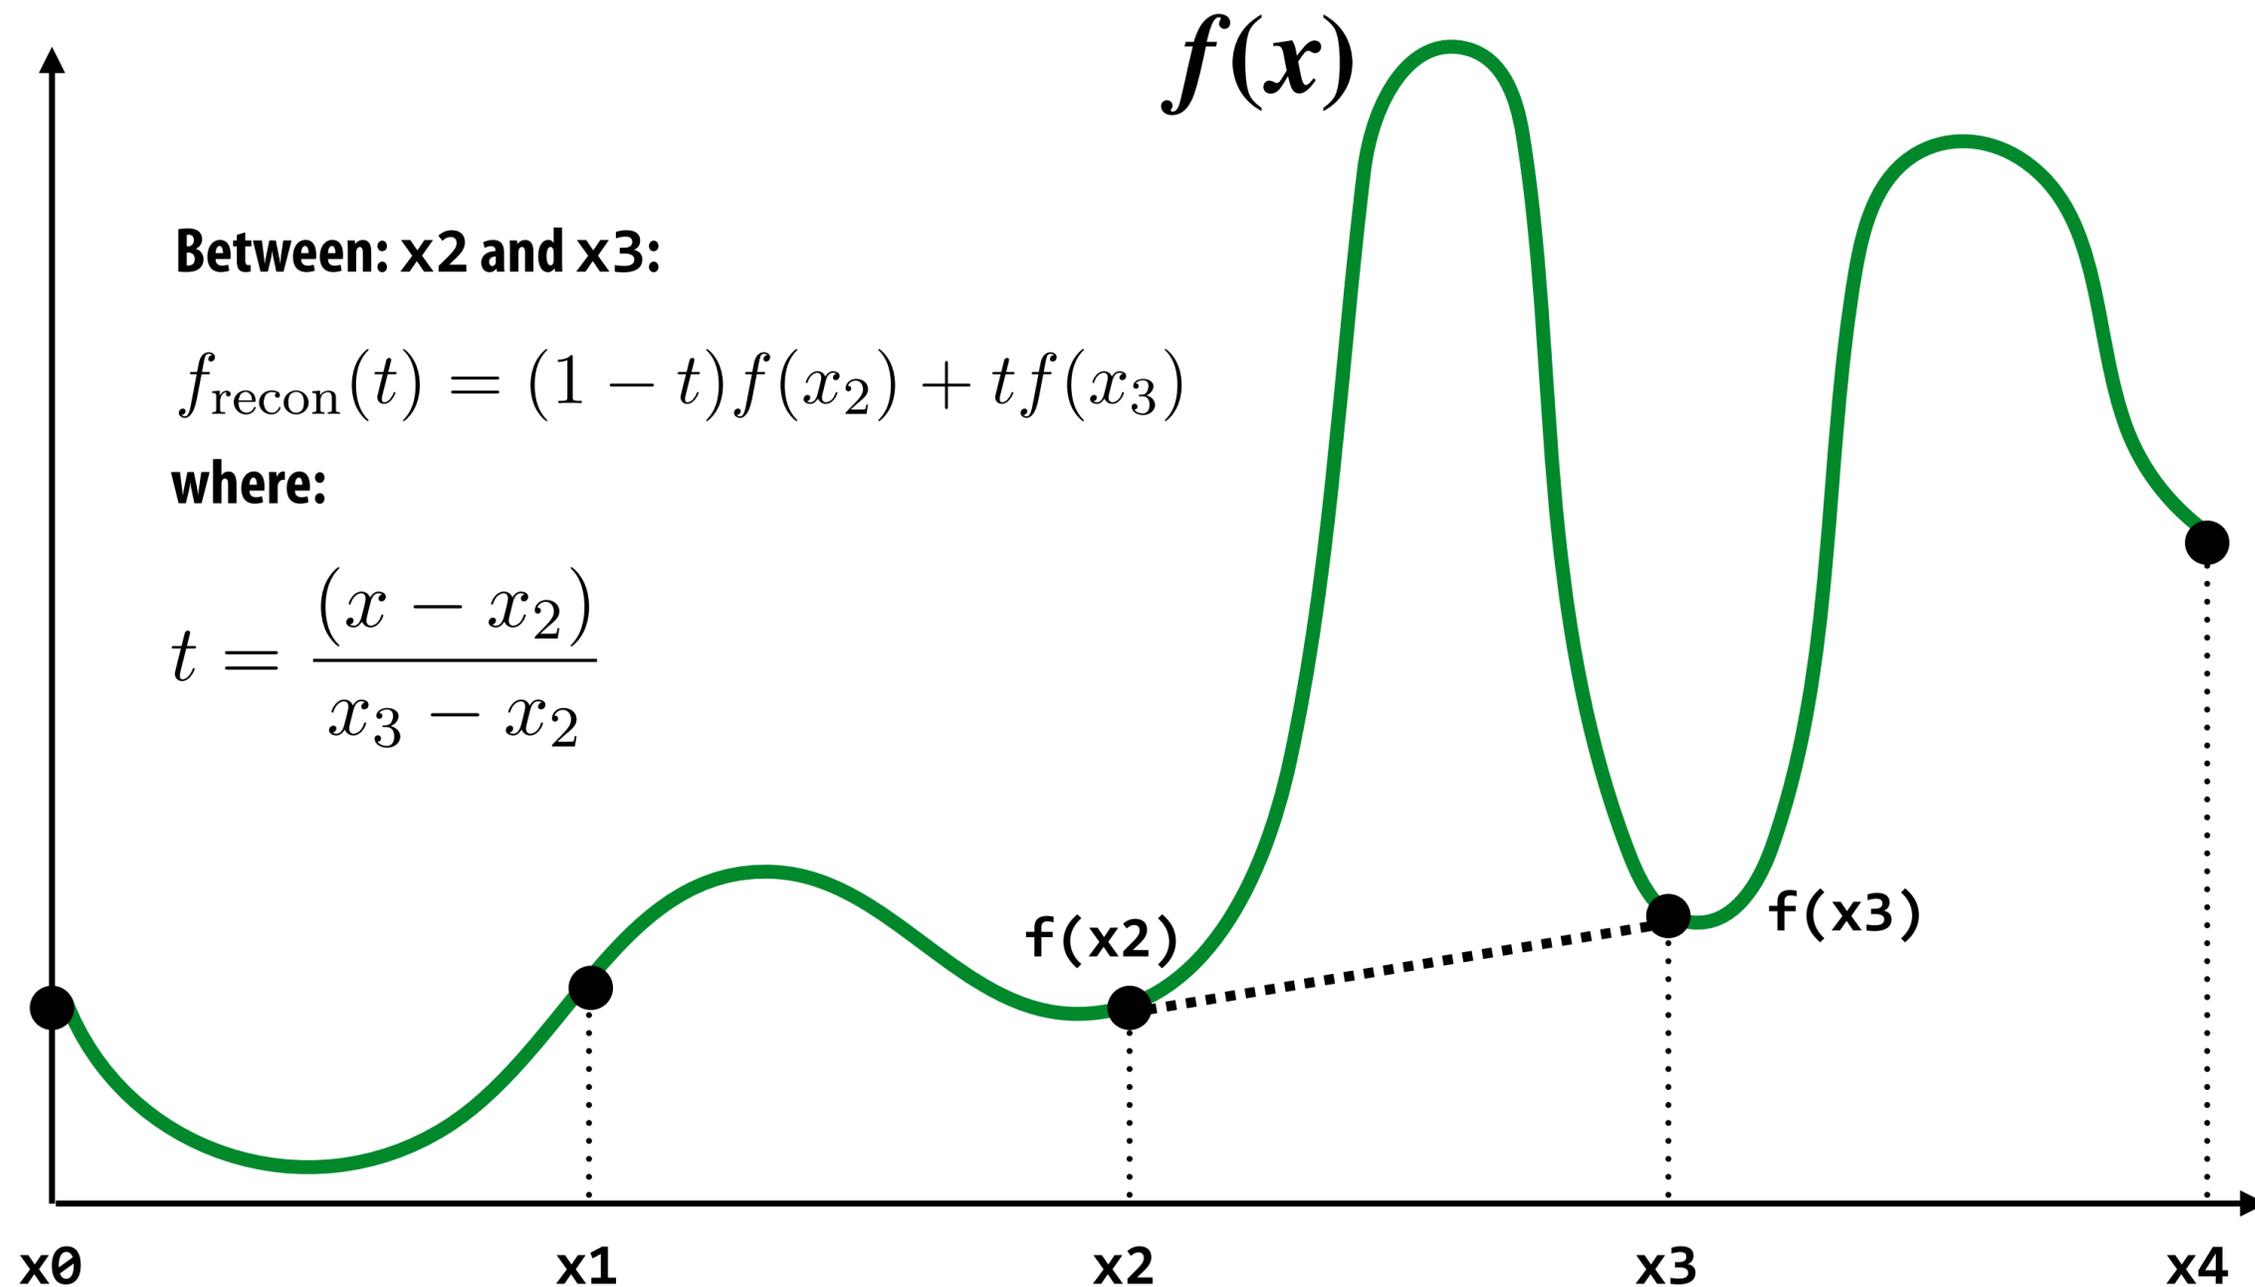
Consider sampling $\text{color}(x,y)$



What is the triangle's color at the point x ?

Review: interpolation in 1D

$f_{\text{recon}}(x)$ = linear interpolation between values of two closest samples to x

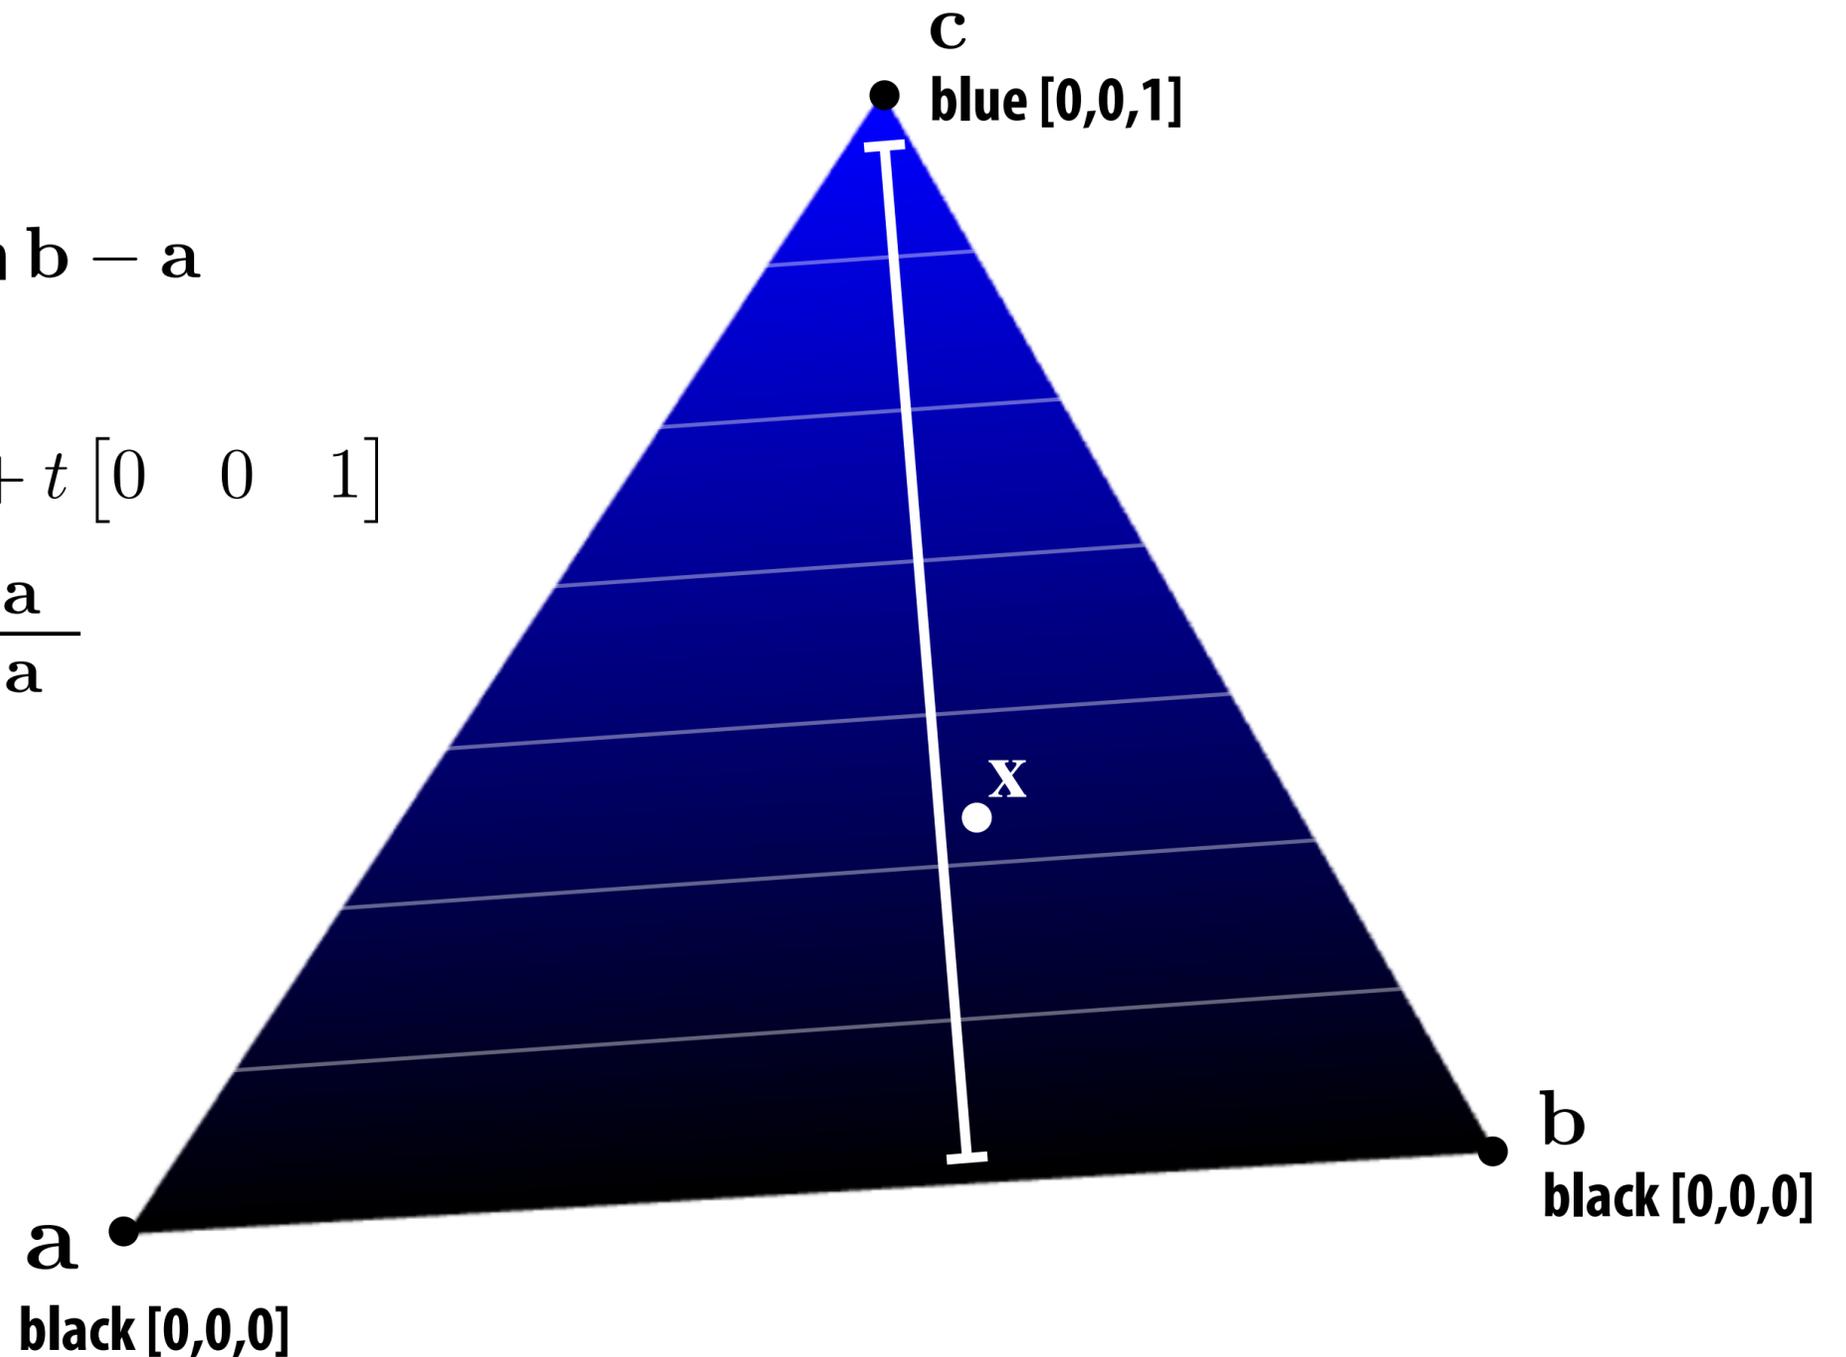


Consider similar behavior on triangle

Color depends on distance from $b - a$

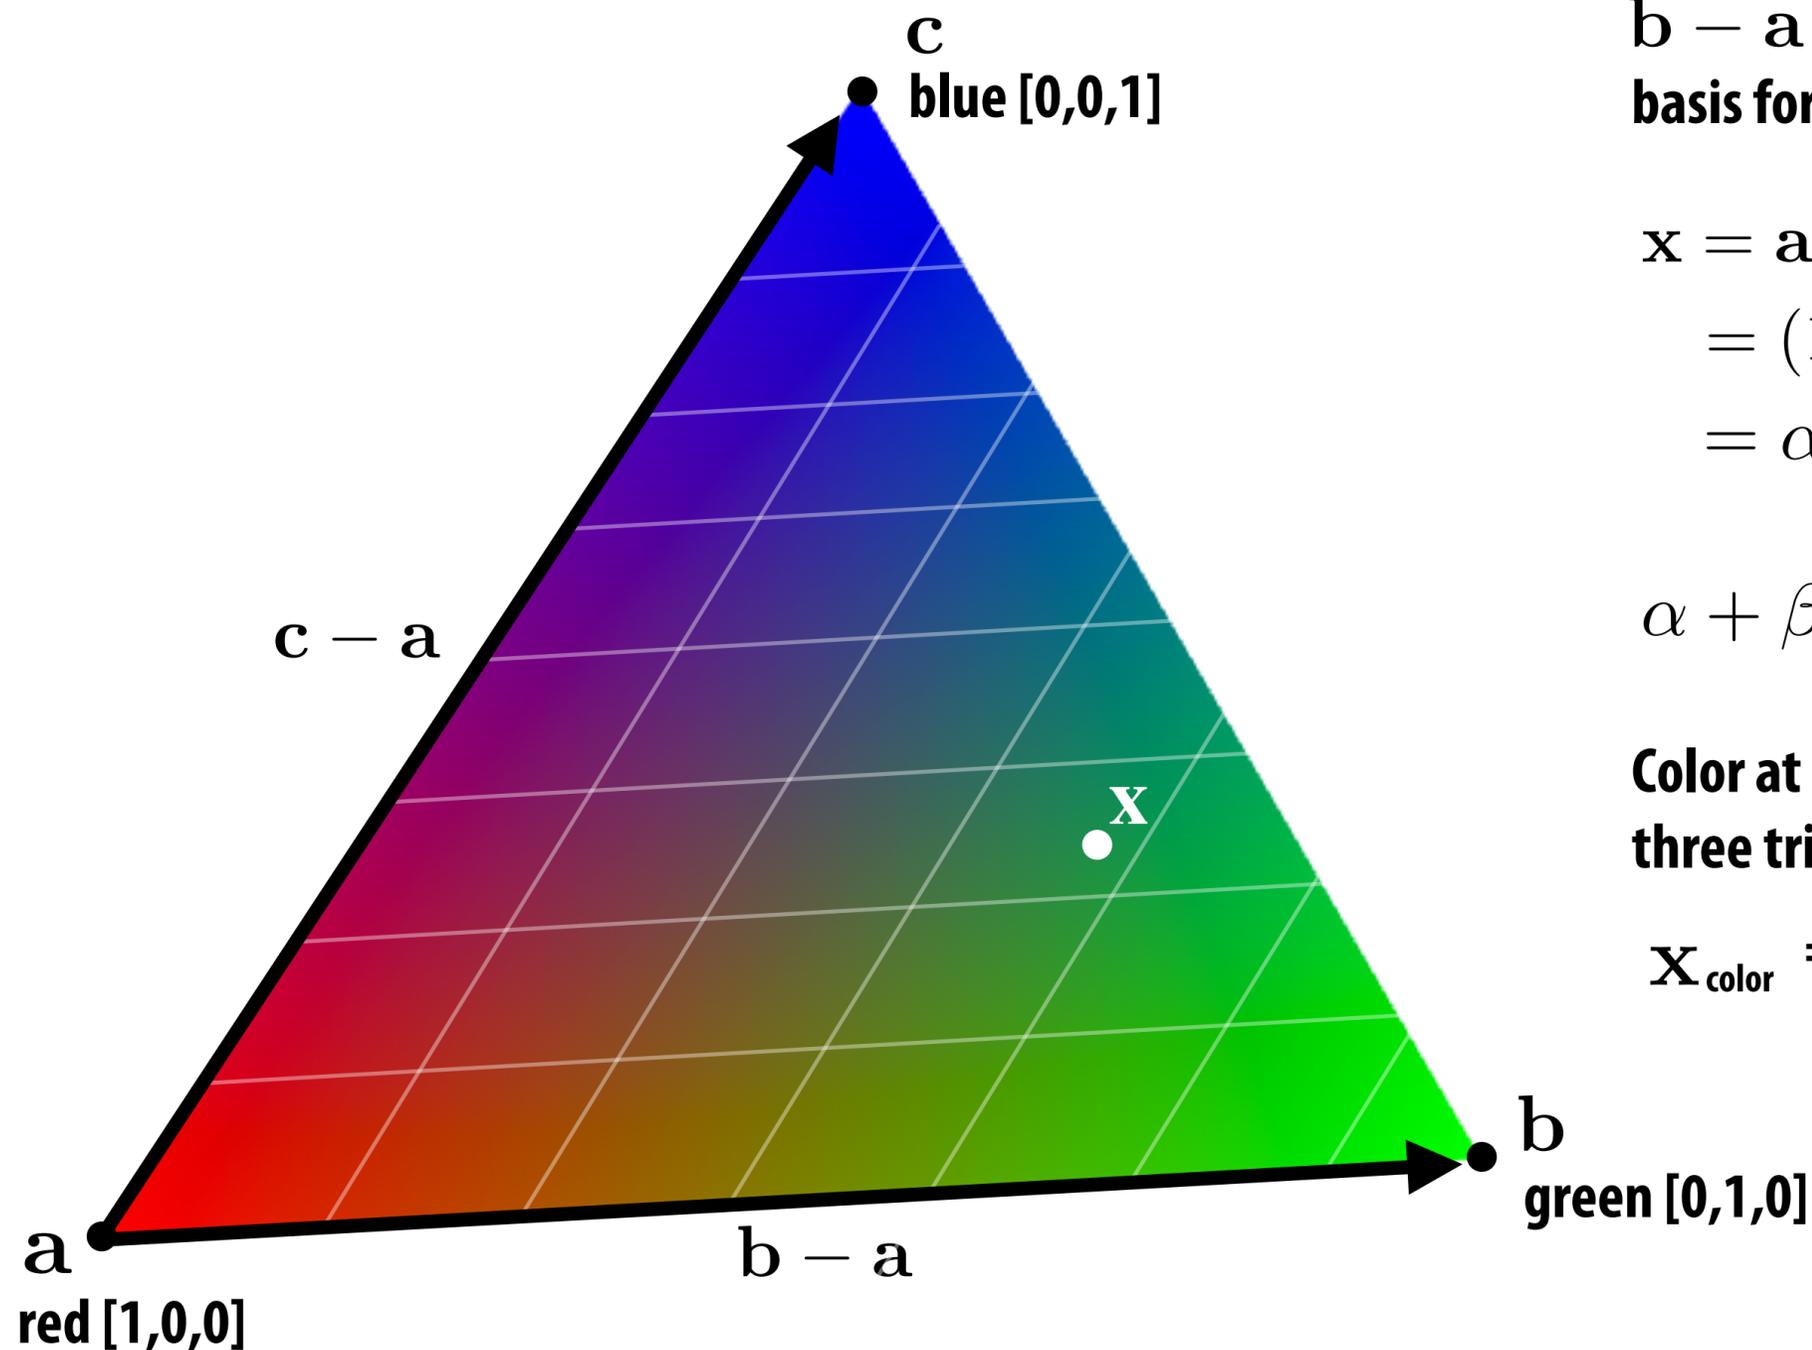
$$\text{color at } (1 - t) \begin{bmatrix} 0 & 0 & 0 \end{bmatrix} + t \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}$$

$$t = \frac{\text{distance from } x \text{ to } b - a}{\text{distance from } c \text{ to } b - a}$$



How can we interpolate in 2D between three values?

Interpolation via barycentric coordinates



$b - a$ and $c - a$ form a non-orthogonal basis for points in triangle (origin at a)

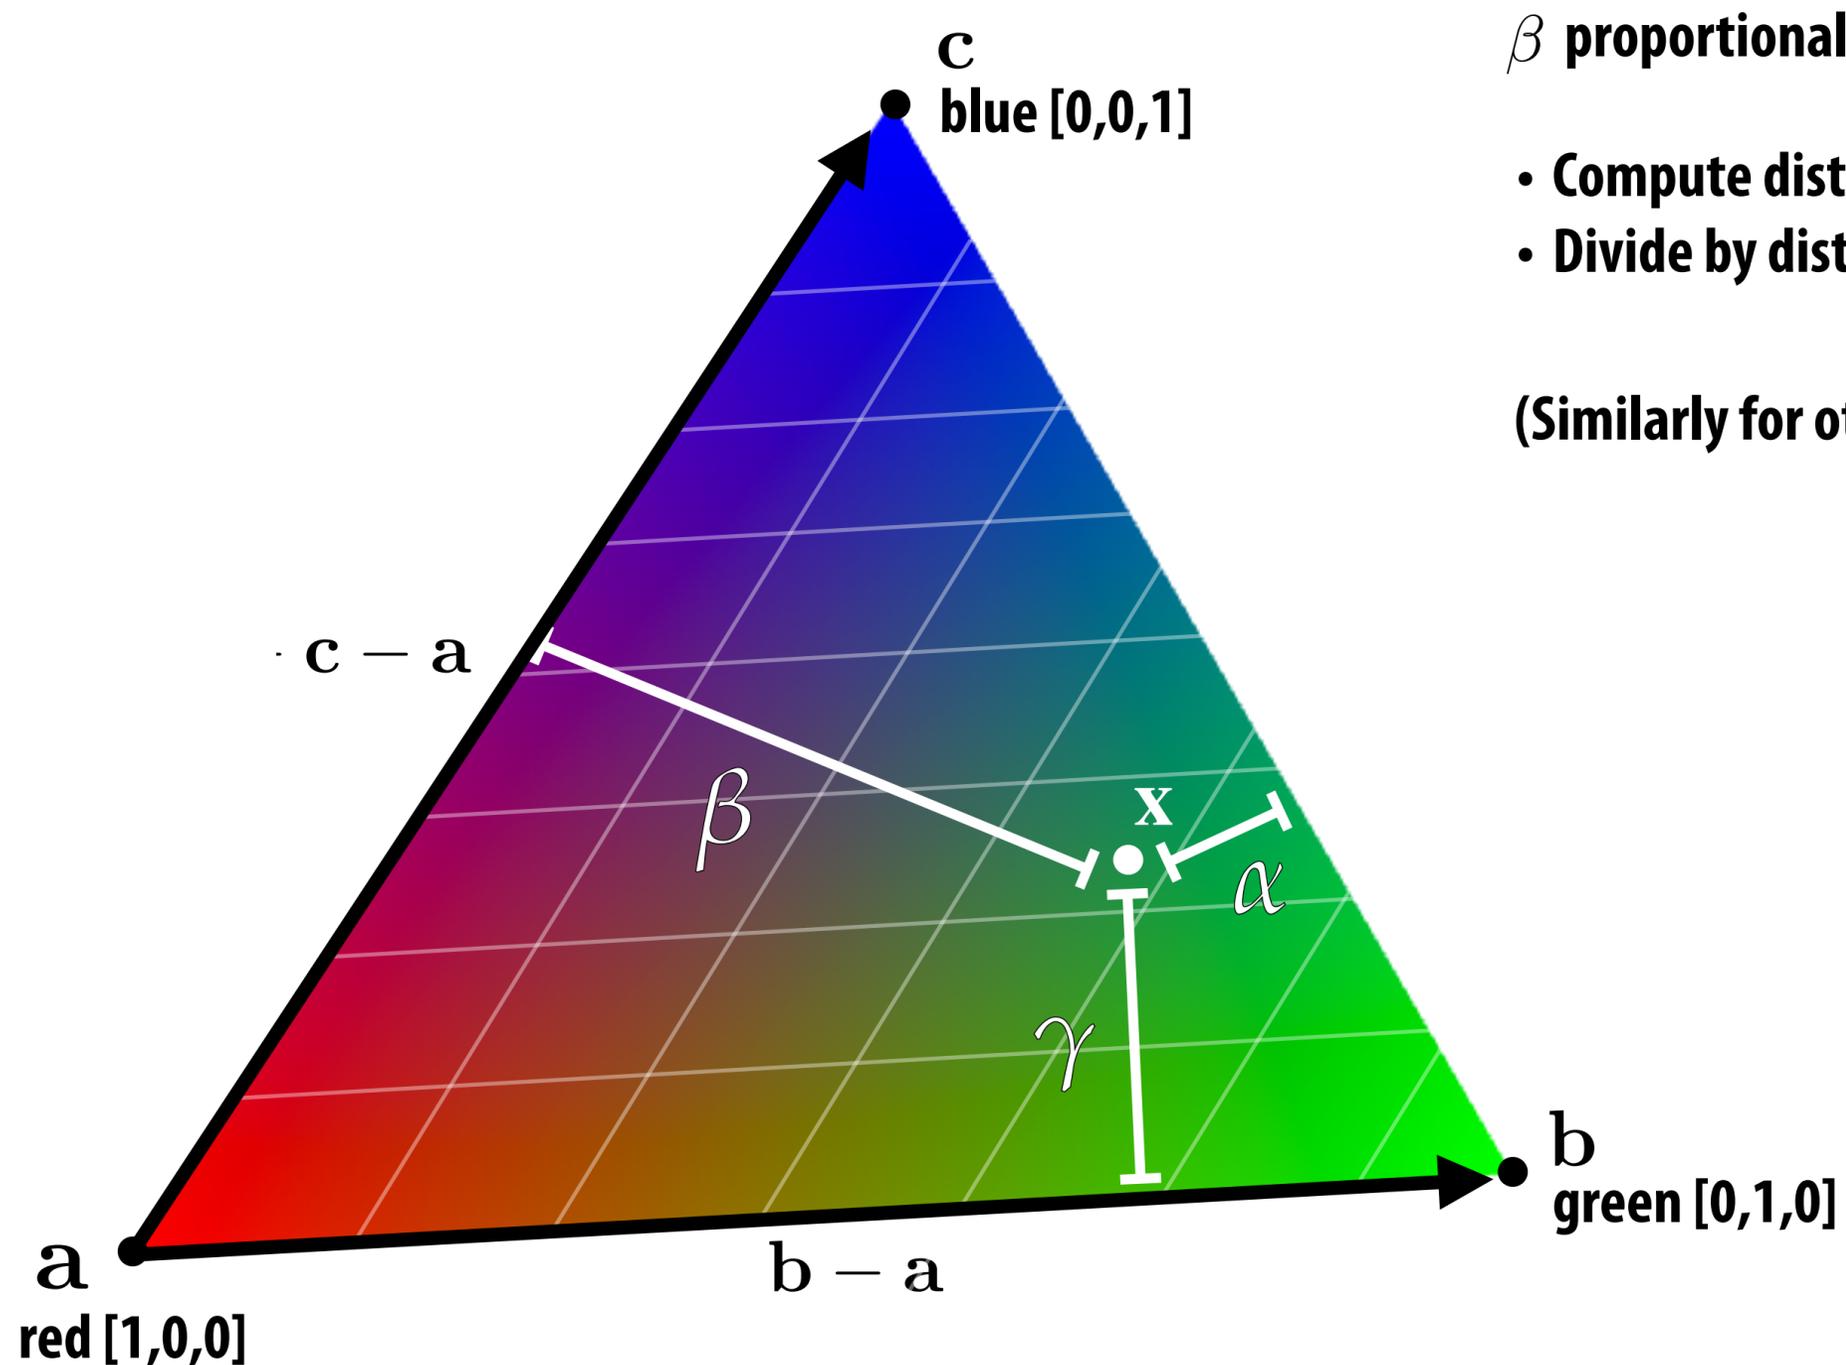
$$\begin{aligned} \mathbf{x} &= \mathbf{a} + \beta(\mathbf{b} - \mathbf{a}) + \gamma(\mathbf{c} - \mathbf{a}) \\ &= (1 - \beta - \gamma)\mathbf{a} + \beta\mathbf{b} + \gamma\mathbf{c} \\ &= \alpha\mathbf{a} + \beta\mathbf{b} + \gamma\mathbf{c} \end{aligned}$$

$$\alpha + \beta + \gamma = 1$$

Color at x is linear combination of color at three triangle vertices.

$$\mathbf{x}_{\text{color}} = \alpha\mathbf{a}_{\text{color}} + \beta\mathbf{b}_{\text{color}} + \gamma\mathbf{c}_{\text{color}}$$

Barycentric coordinates as scaled distances

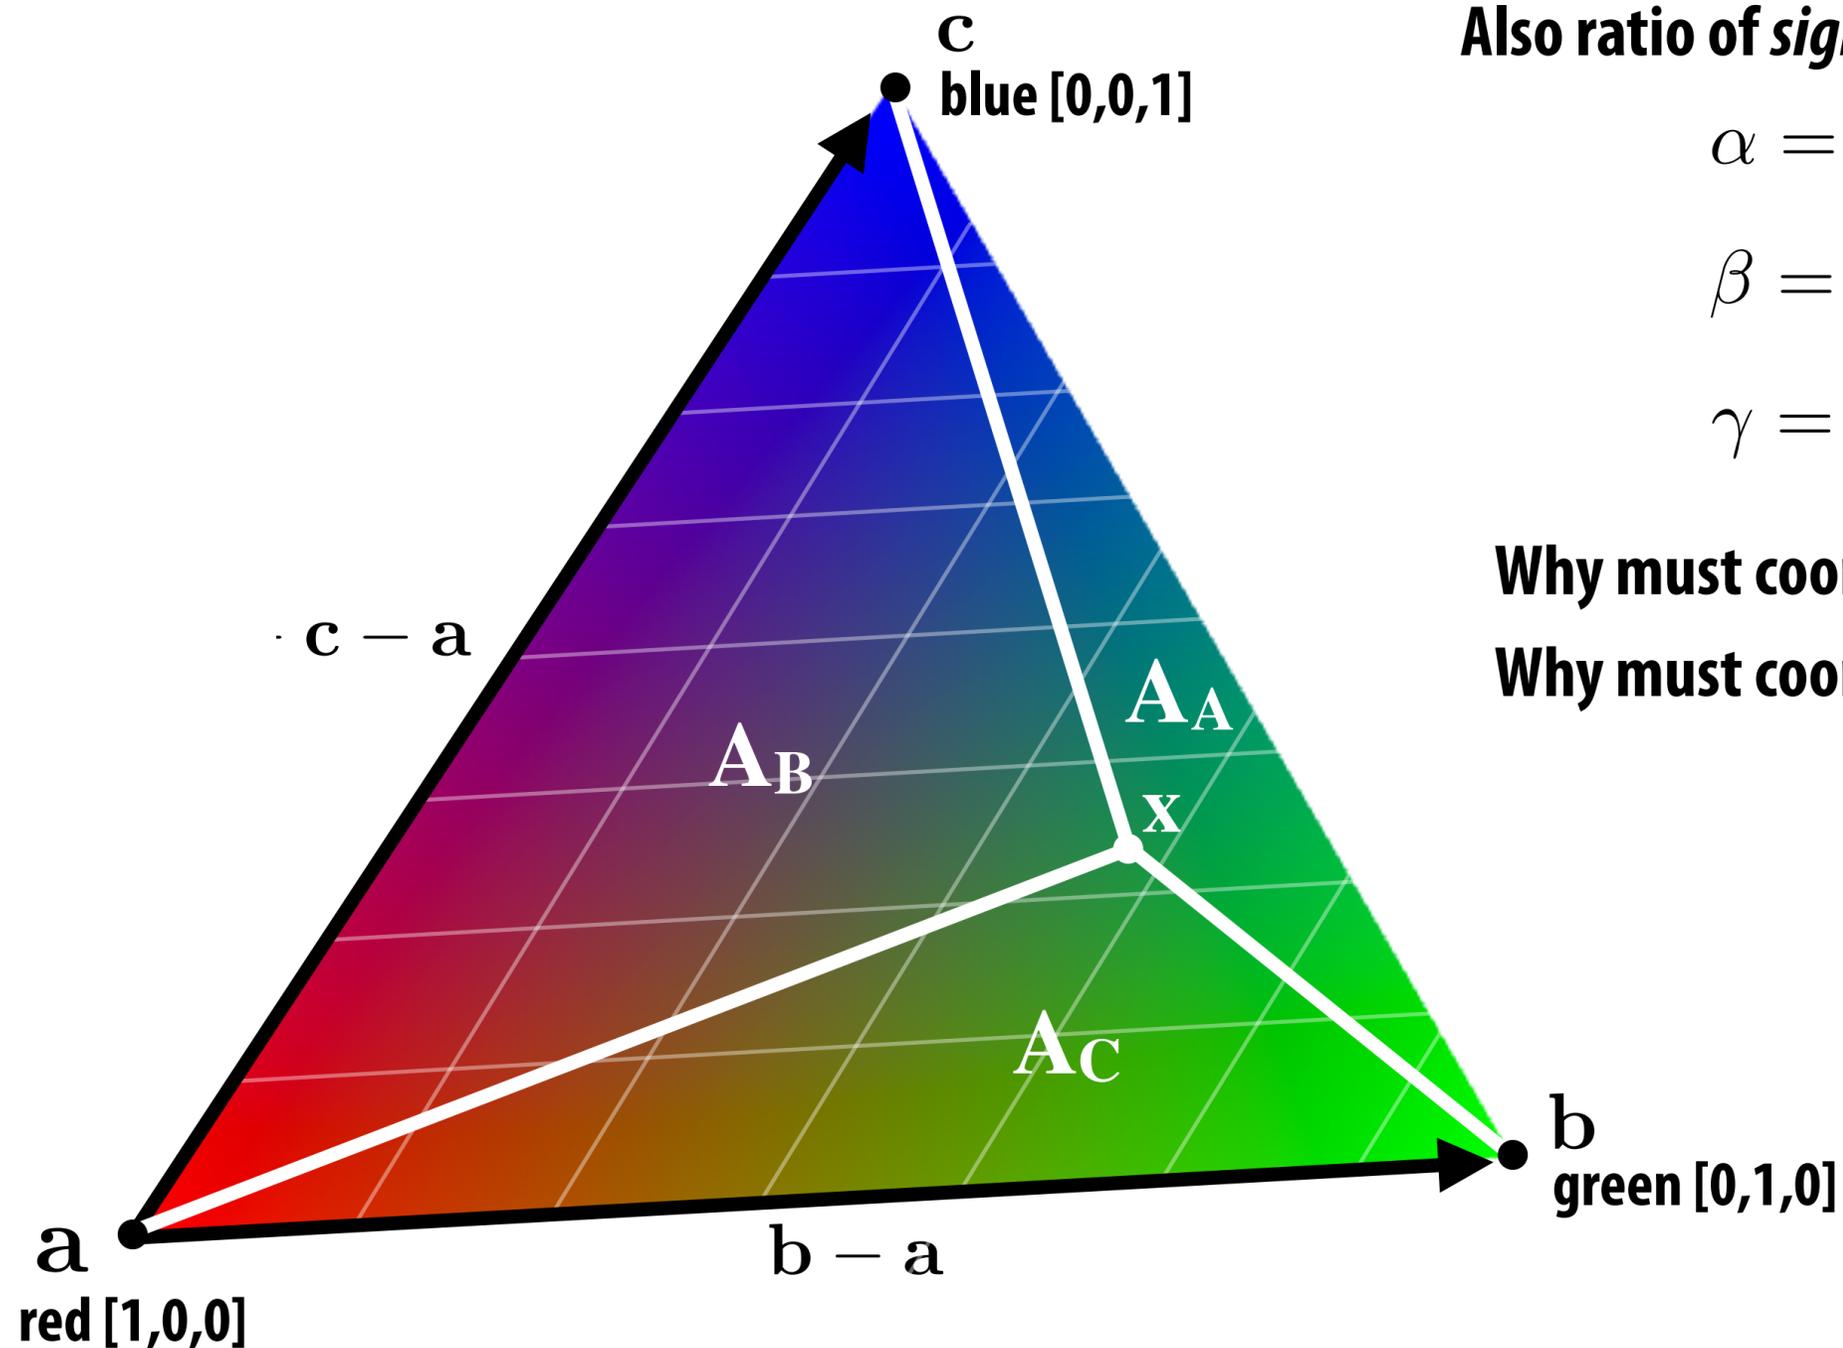


β proportional to distance from x to edge $c - a$

- Compute distance of x from line ca
- Divide by distance of b from line ca ("height")

(Similarly for other two barycentric coordinates)

Barycentric coordinates as ratio of areas



Also ratio of *signed* areas:

$$\alpha = A_A/A$$

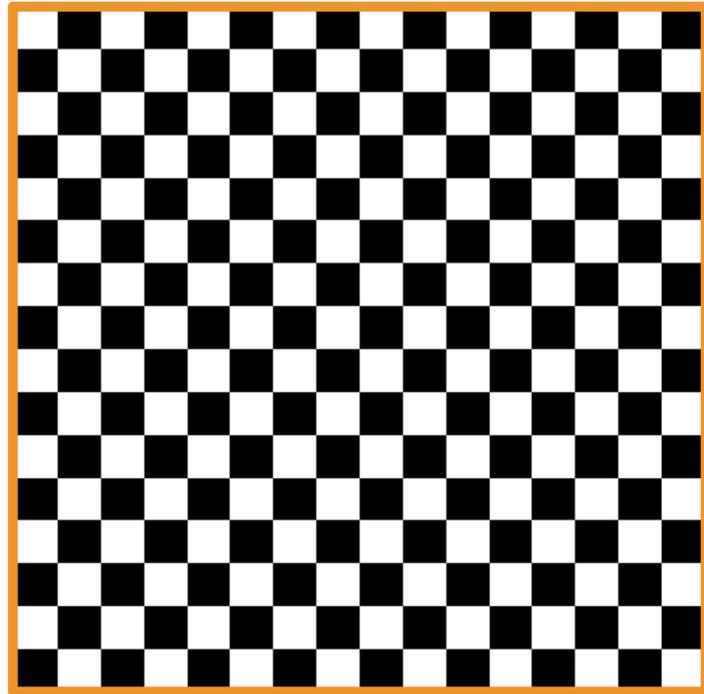
$$\beta = A_B/A$$

$$\gamma = A_C/A$$

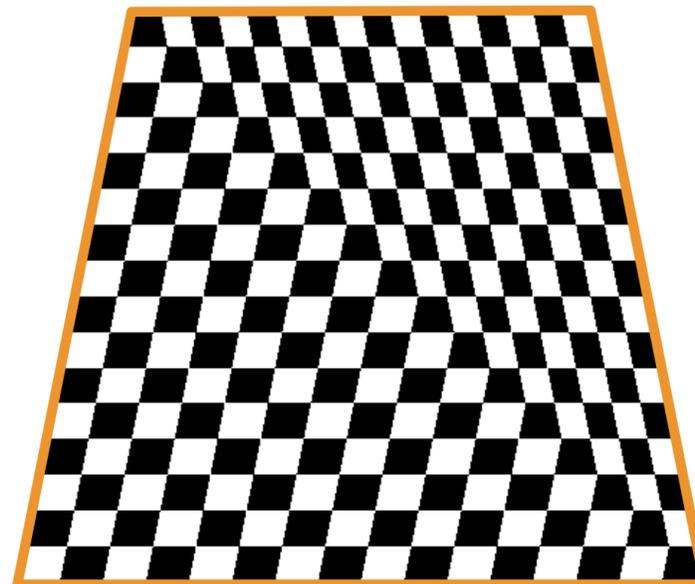
Why must coordinates sum to one?

Why must coordinates be between 0 and 1?

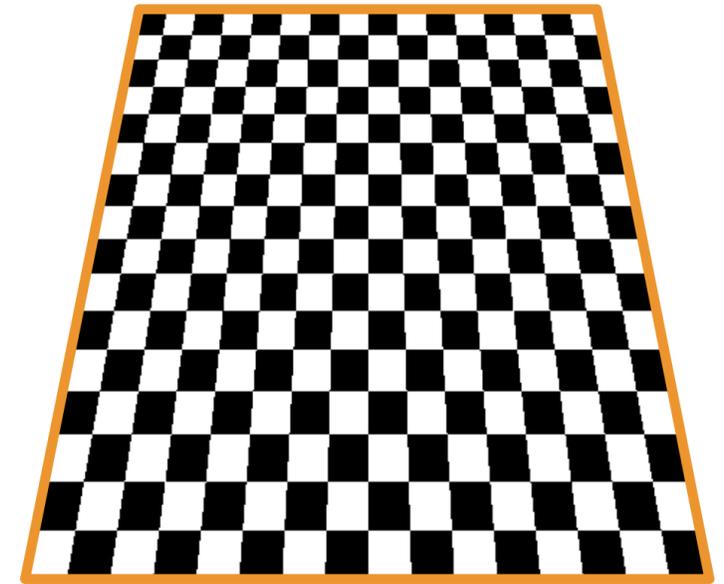
Perspective projection and interpolation



Texture



Plane tilted
down with
perspective
projection —
What's wrong?

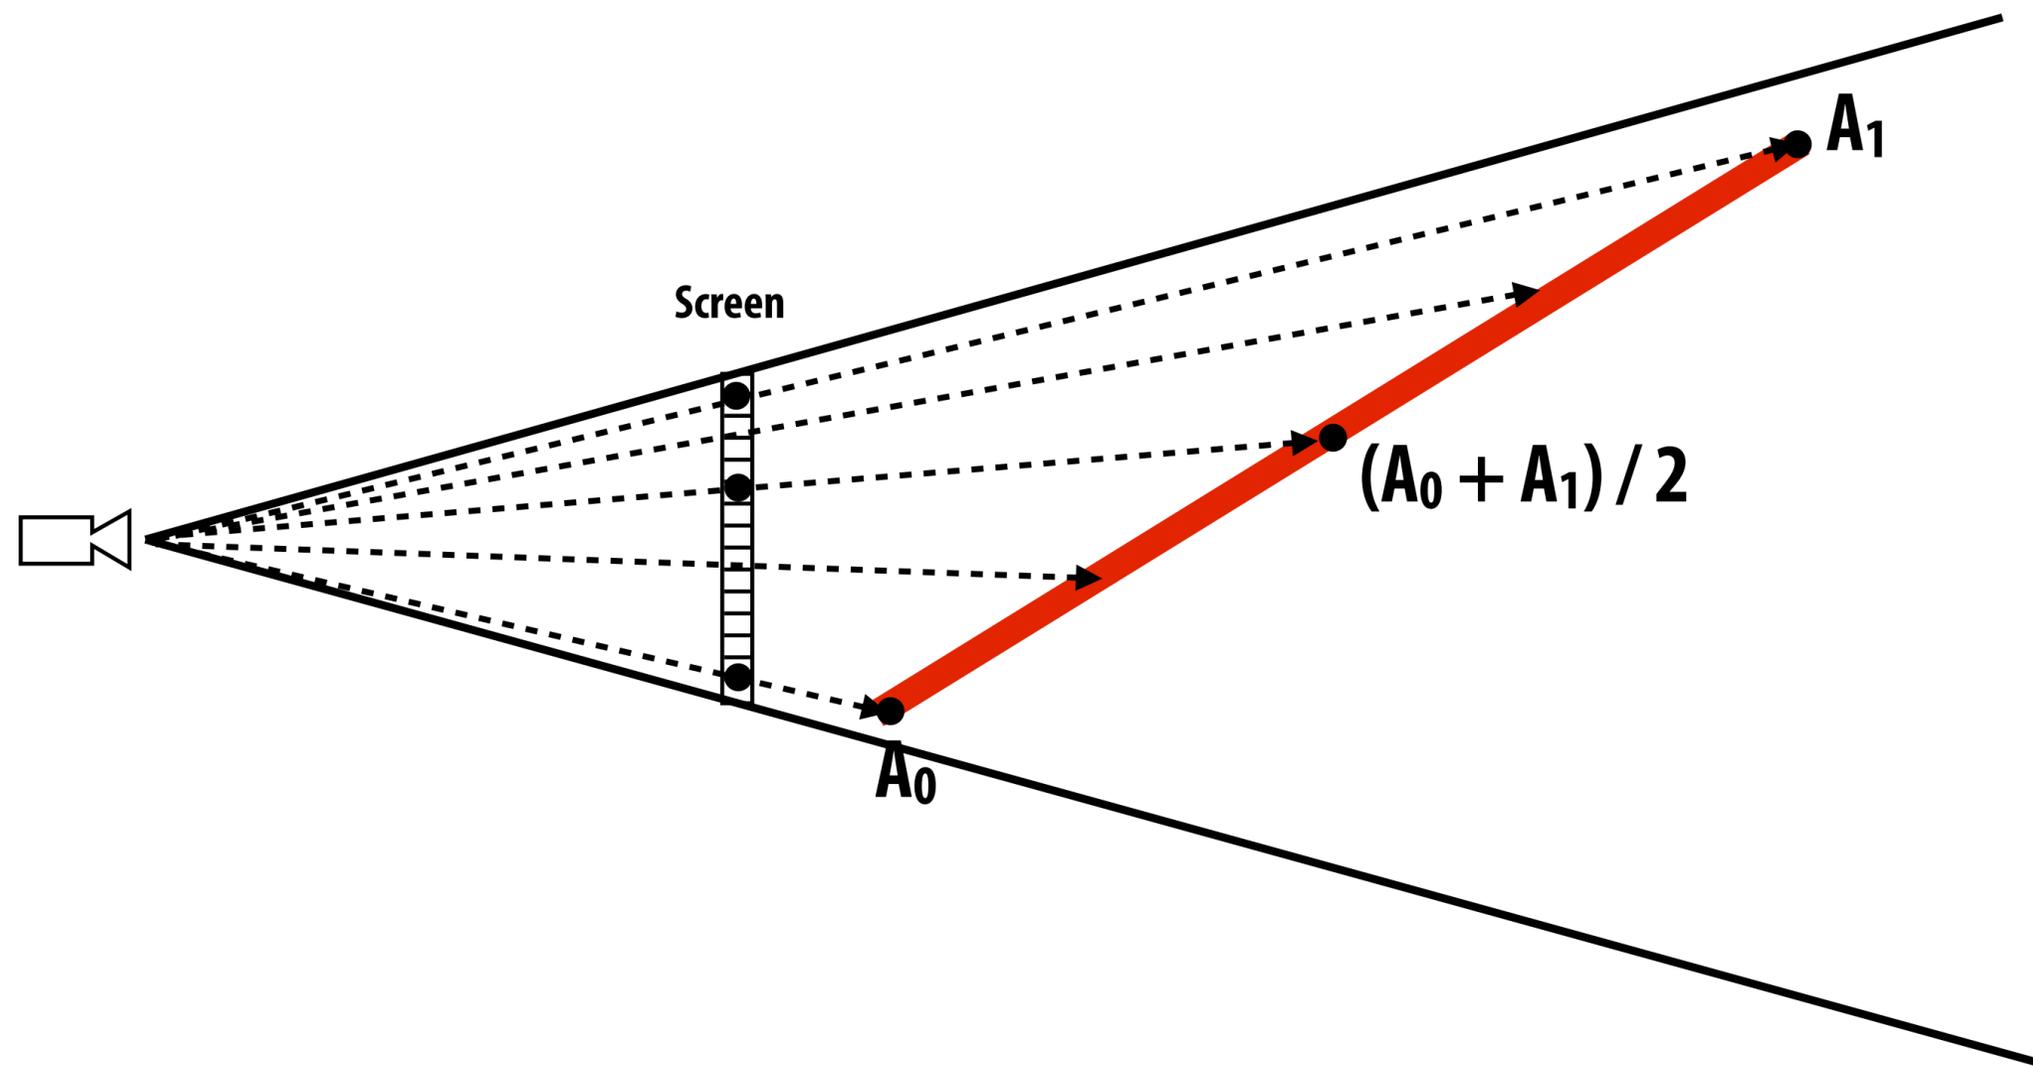


Correct image

Perspective incorrect interpolation

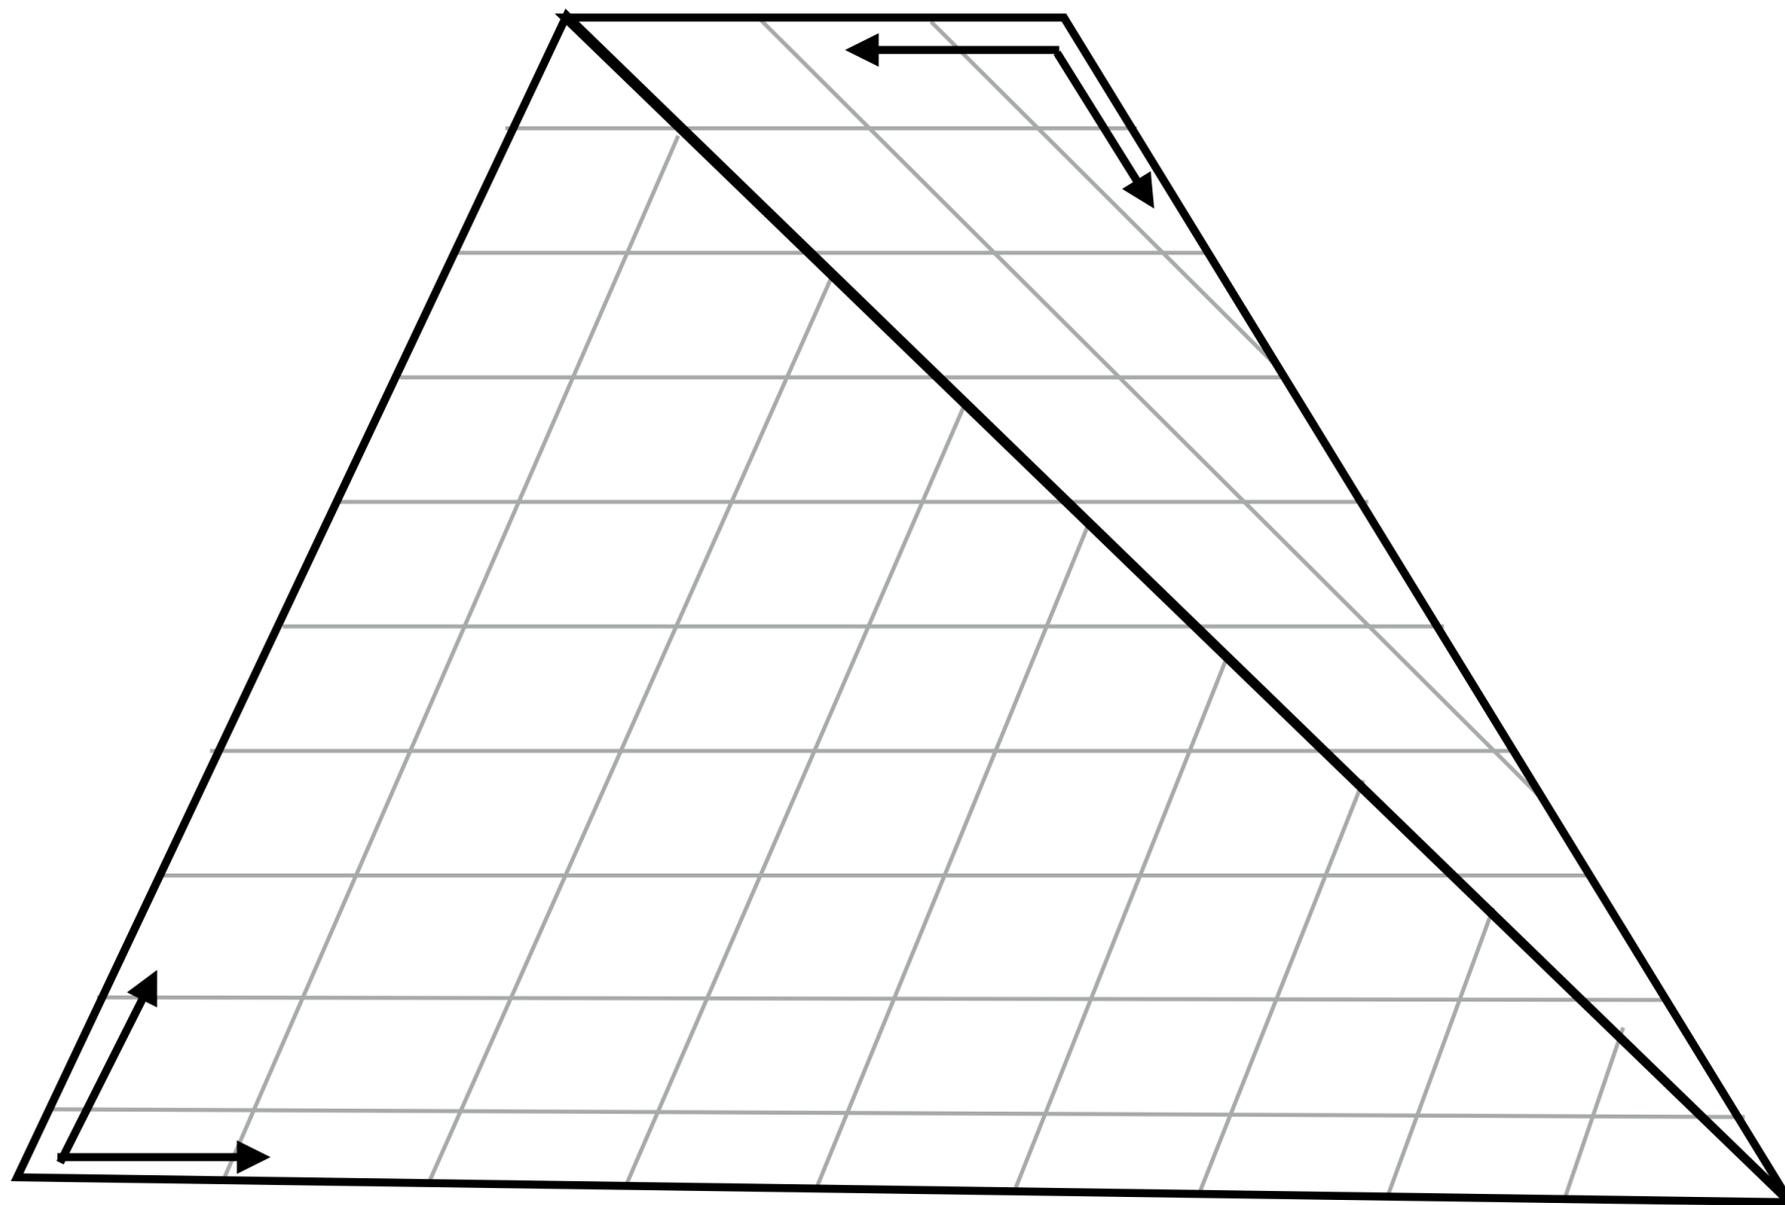
Due to perspective projection (homogeneous divide), barycentric interpolation of values on a triangle with different depths is not an affine function of screen XY coordinates.

Attribute values must be interpolated linearly in 3D space.



Example: perspective incorrect interpolation

Good example is quadrilateral split into two triangles:



If we compute barycentric coordinates using 2D (projected) coordinates, can lead to (derivative) discontinuity in interpolation where quad was split.

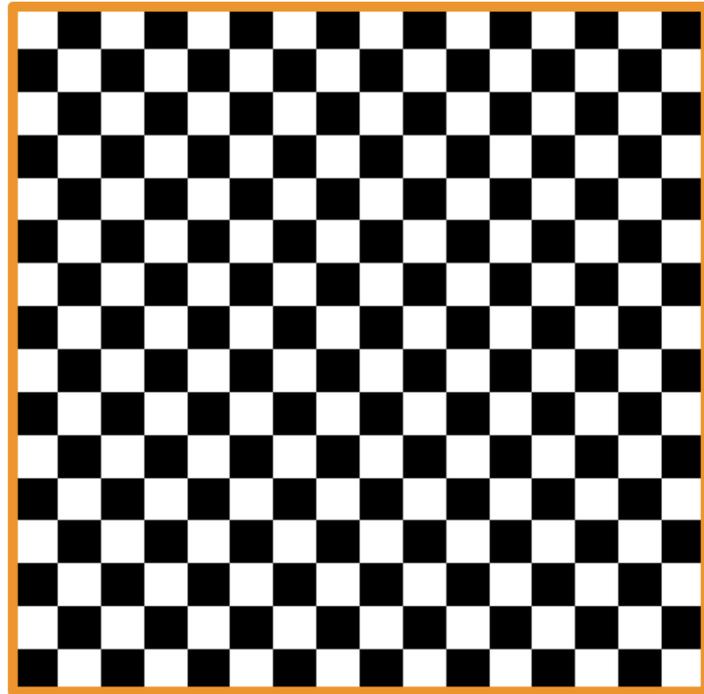
Perspective correct interpolation

■ Basic recipe:

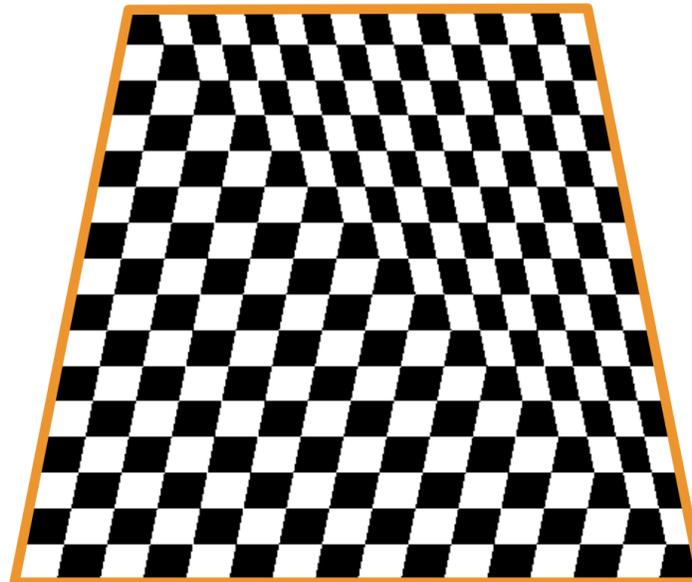
- To interpolate some attribute ϕ ...
- Compute depth z at each vertex
- Evaluate $Z := 1/z$ and $P := \phi/z$ at each vertex
- Interpolate Z and P using standard (2D) barycentric coords
- At each *fragment*, divide interpolated P by interpolated Z to get final value of ϕ



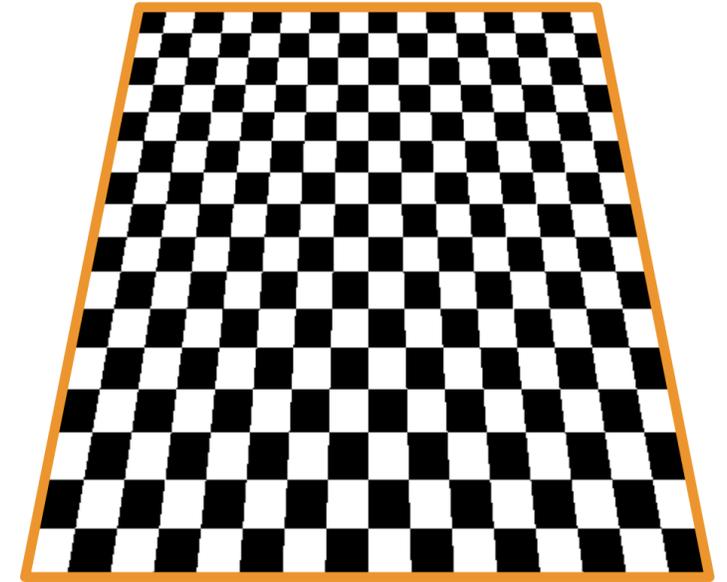
Perspective correct interpolation



Texture



**Affine
screen-space
interpolation**



**Perspective
world-space
interpolation**

Texture mapping



Many uses of texture mapping

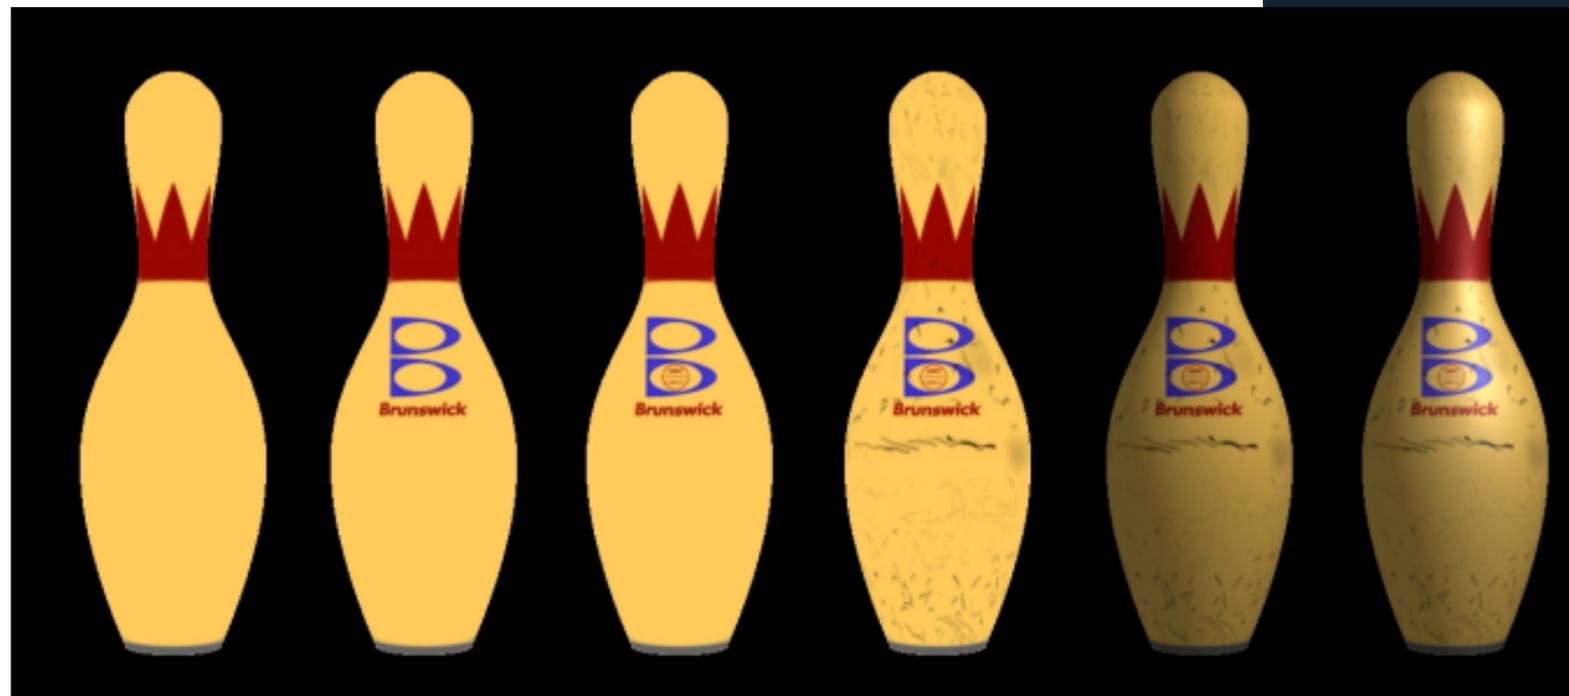
Define variation in surface reflectance



Pattern on ball

Wood grain on floor

Describe surface material properties

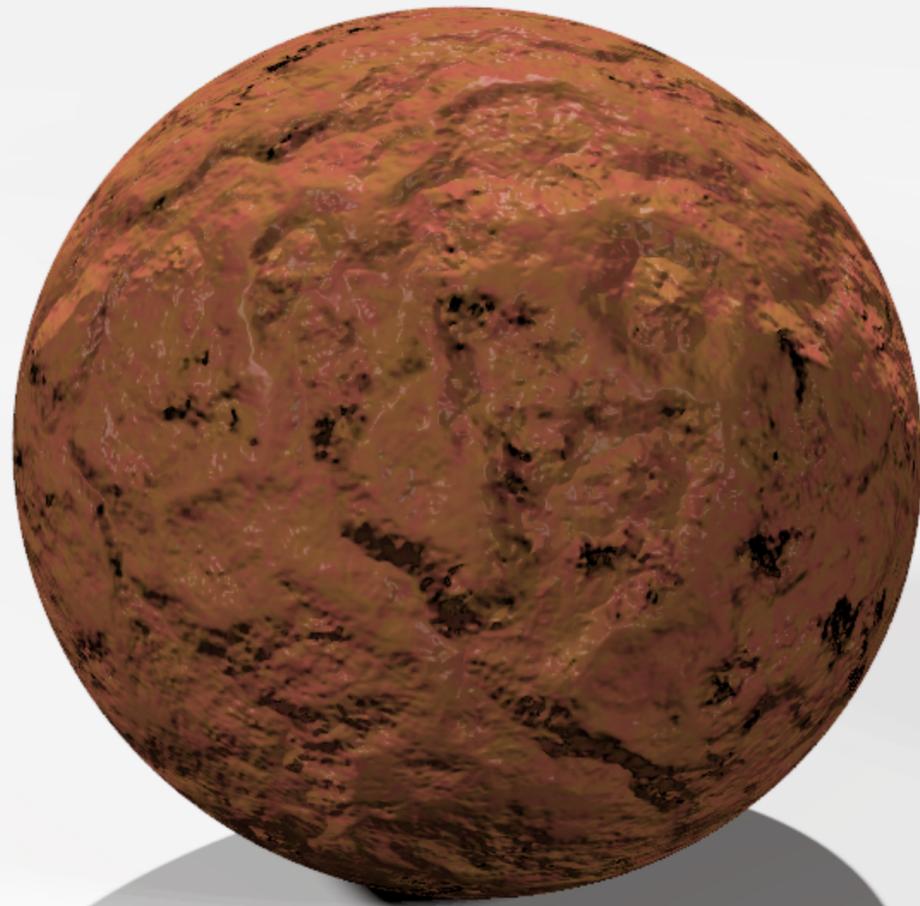


Multiple layers of texture maps for color, logos, scratches, etc.



Normal and displacement mapping

normal mapping



Use texture value to perturb surface normal to “fake” appearance of a bumpy surface (note smooth silhouette/shadow reveals that surface geometry is not actually bumpy!)

displacement mapping



dice up surface geometry into tiny triangles & offset positions according to texture values (note bumpy silhouette and shadow boundary)

Represent precomputed lighting and shadows



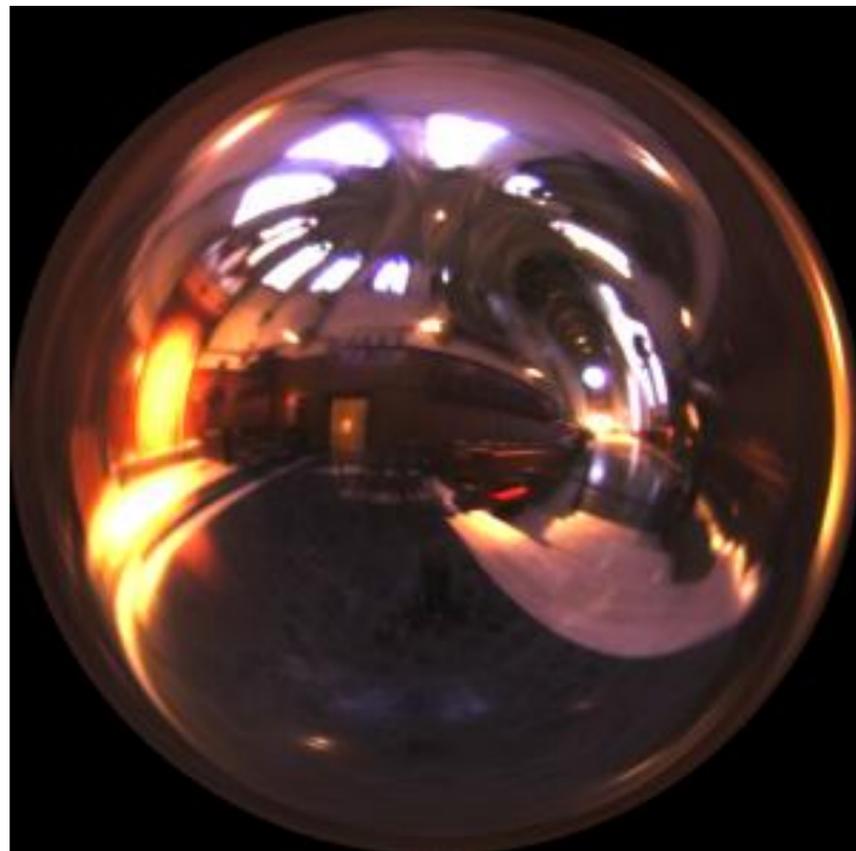
Original model



With ambient occlusion



Extracted ambient occlusion map



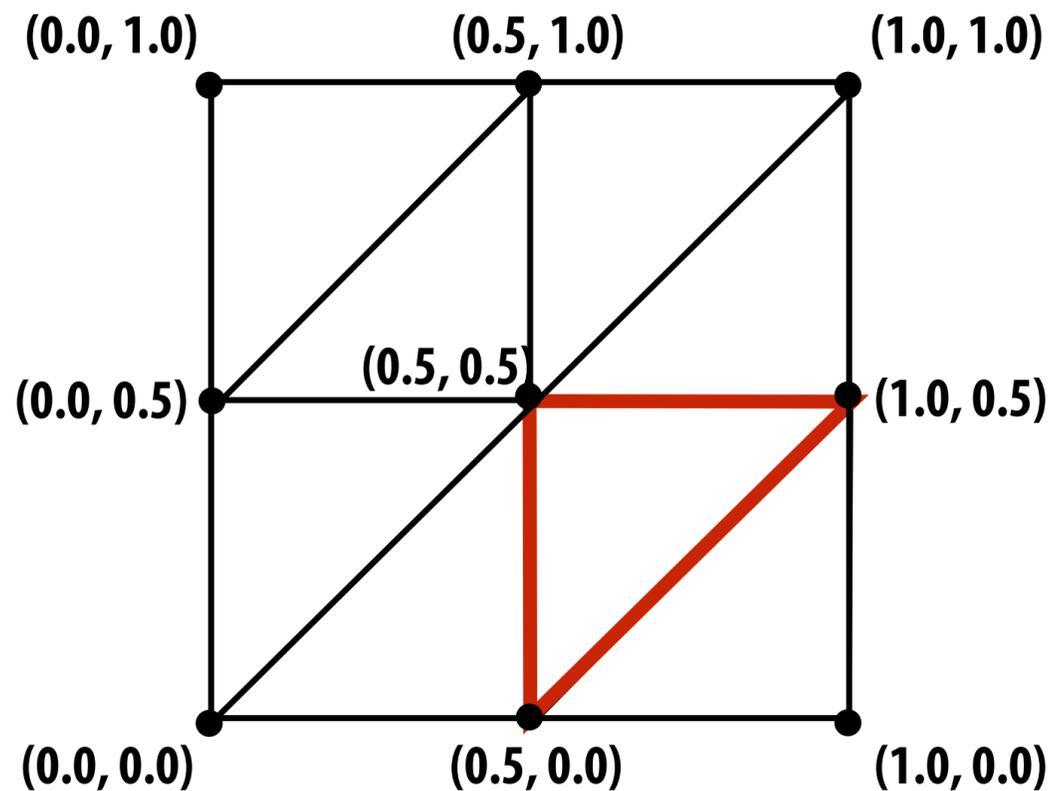
Grace Cathedral environment map



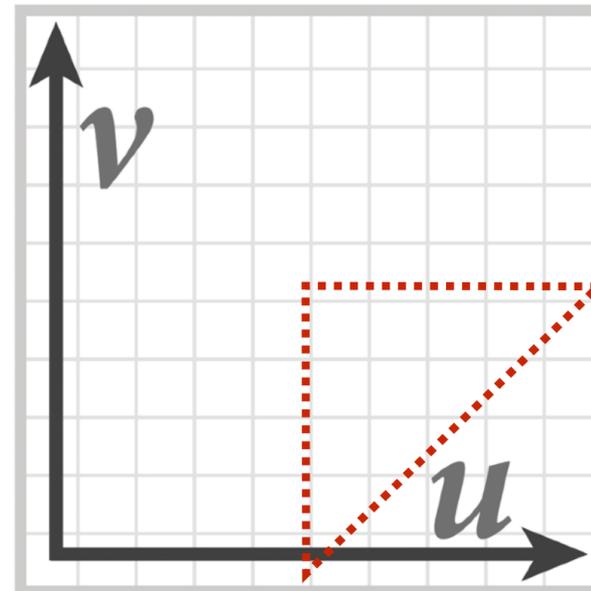
Environment map used in rendering

Texture coordinates

“Texture coordinates” define a mapping from surface coordinates (points on triangle) to points in texture domain.

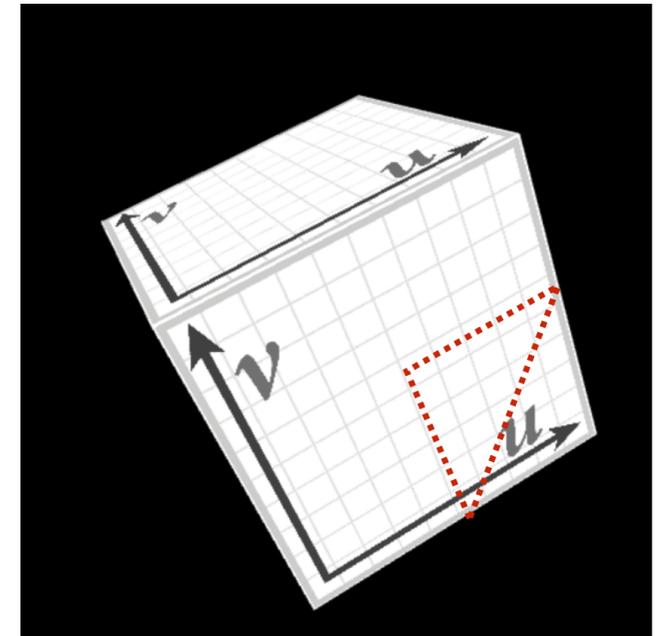


Eight triangles (one face of cube) with surface parameterization provided as per-vertex texture coordinates.



$\text{myTex}(u, v)$ is a function defined on the $[0, 1]^2$ domain (represented by 2048x2048 image)

Location of highlighted triangle in texture space shown in red.



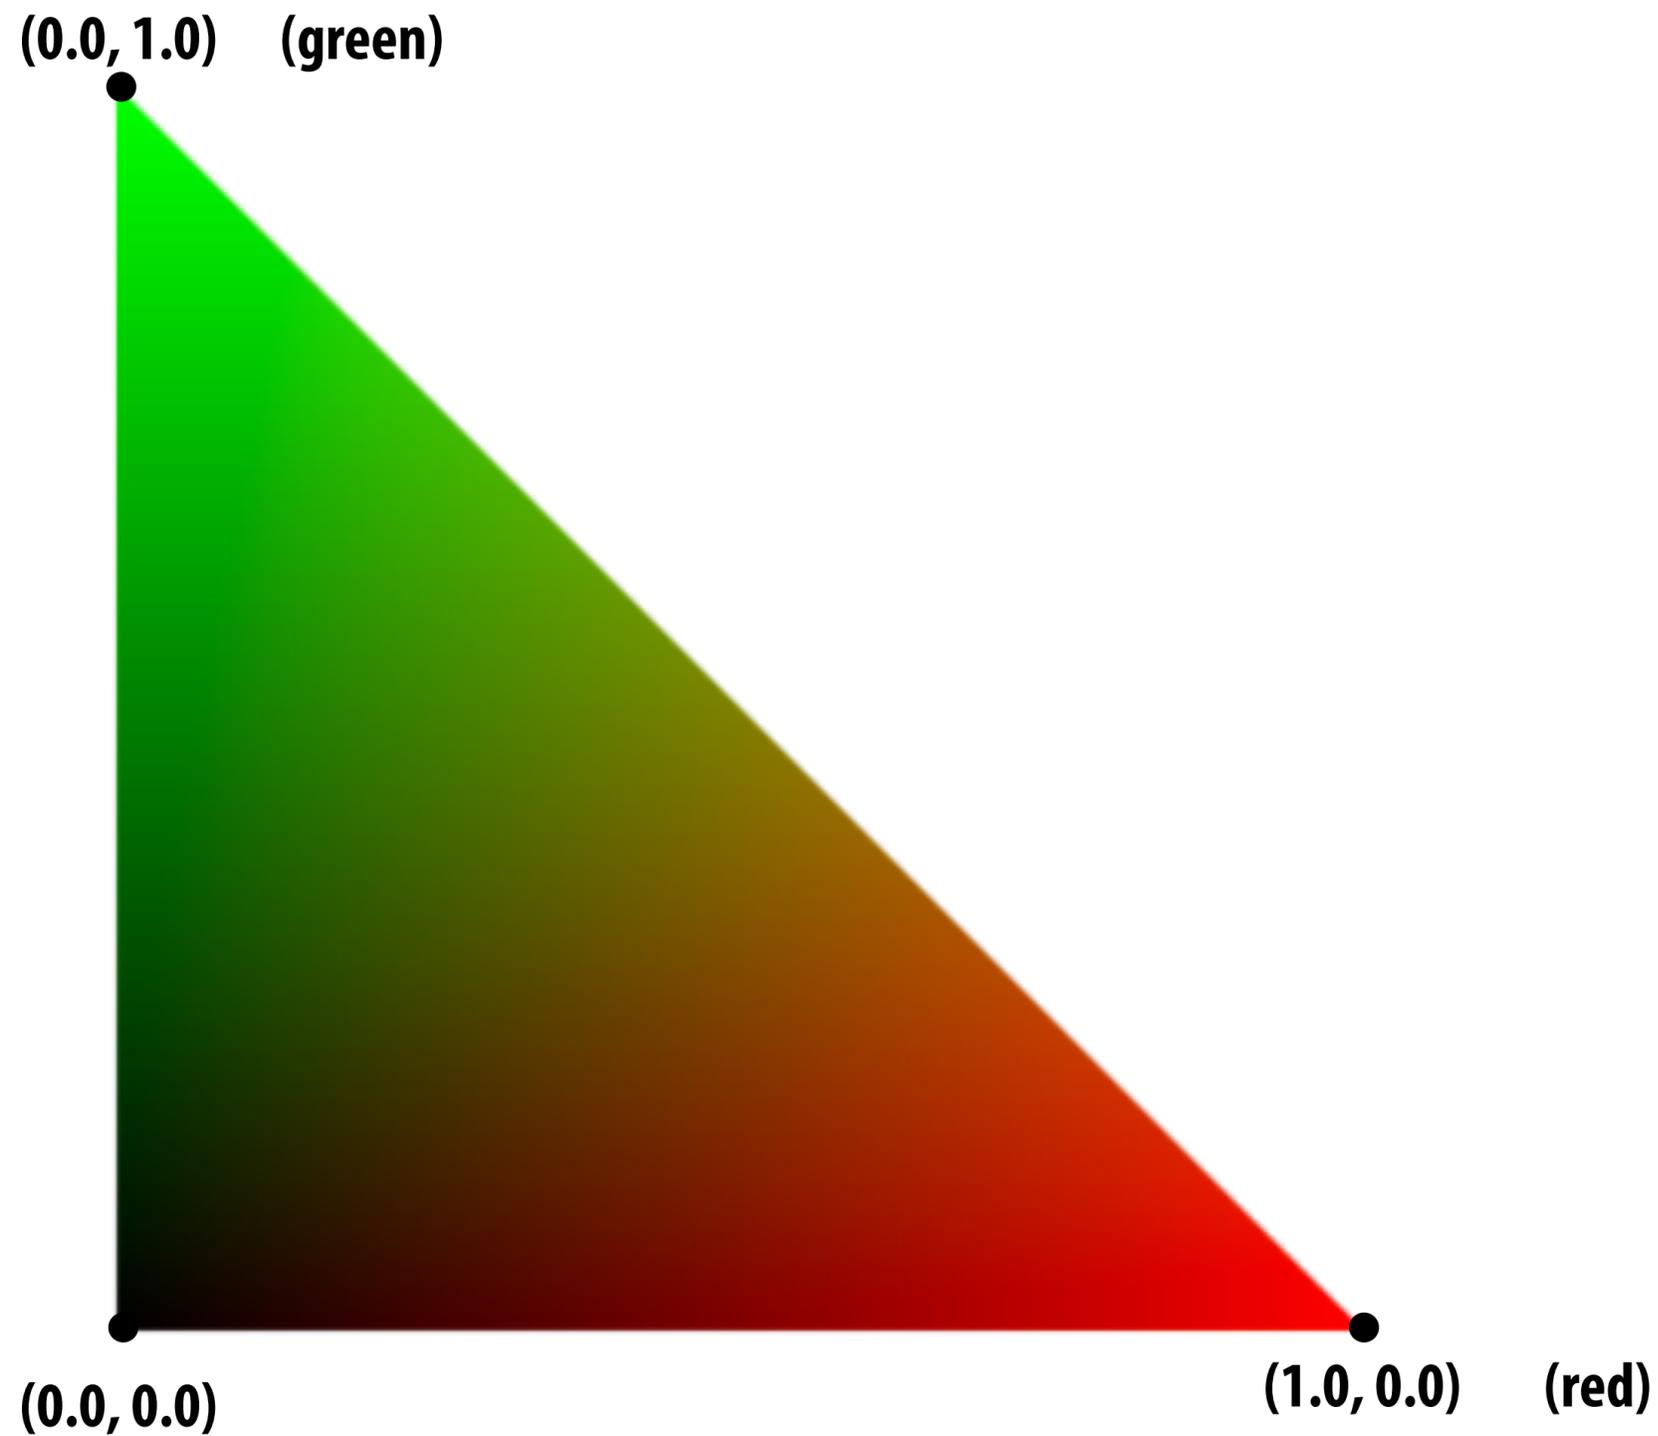
Final rendered result (entire cube shown).

Location of triangle after projection onto screen shown in red.

(We'll assume surface-to-texture space mapping is provided as per vertex values)

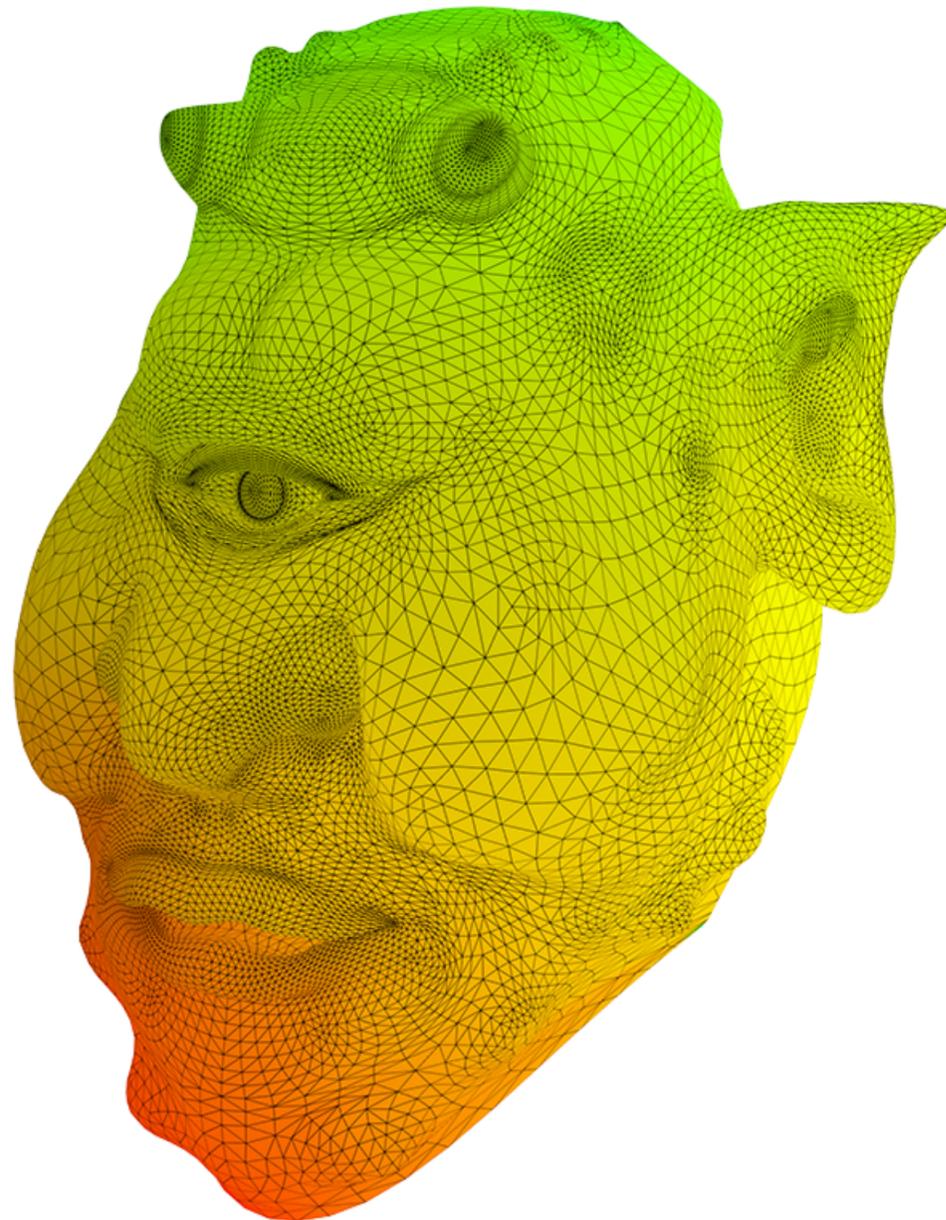
Visualization of texture coordinates

Texture coordinates linearly interpolated over triangle

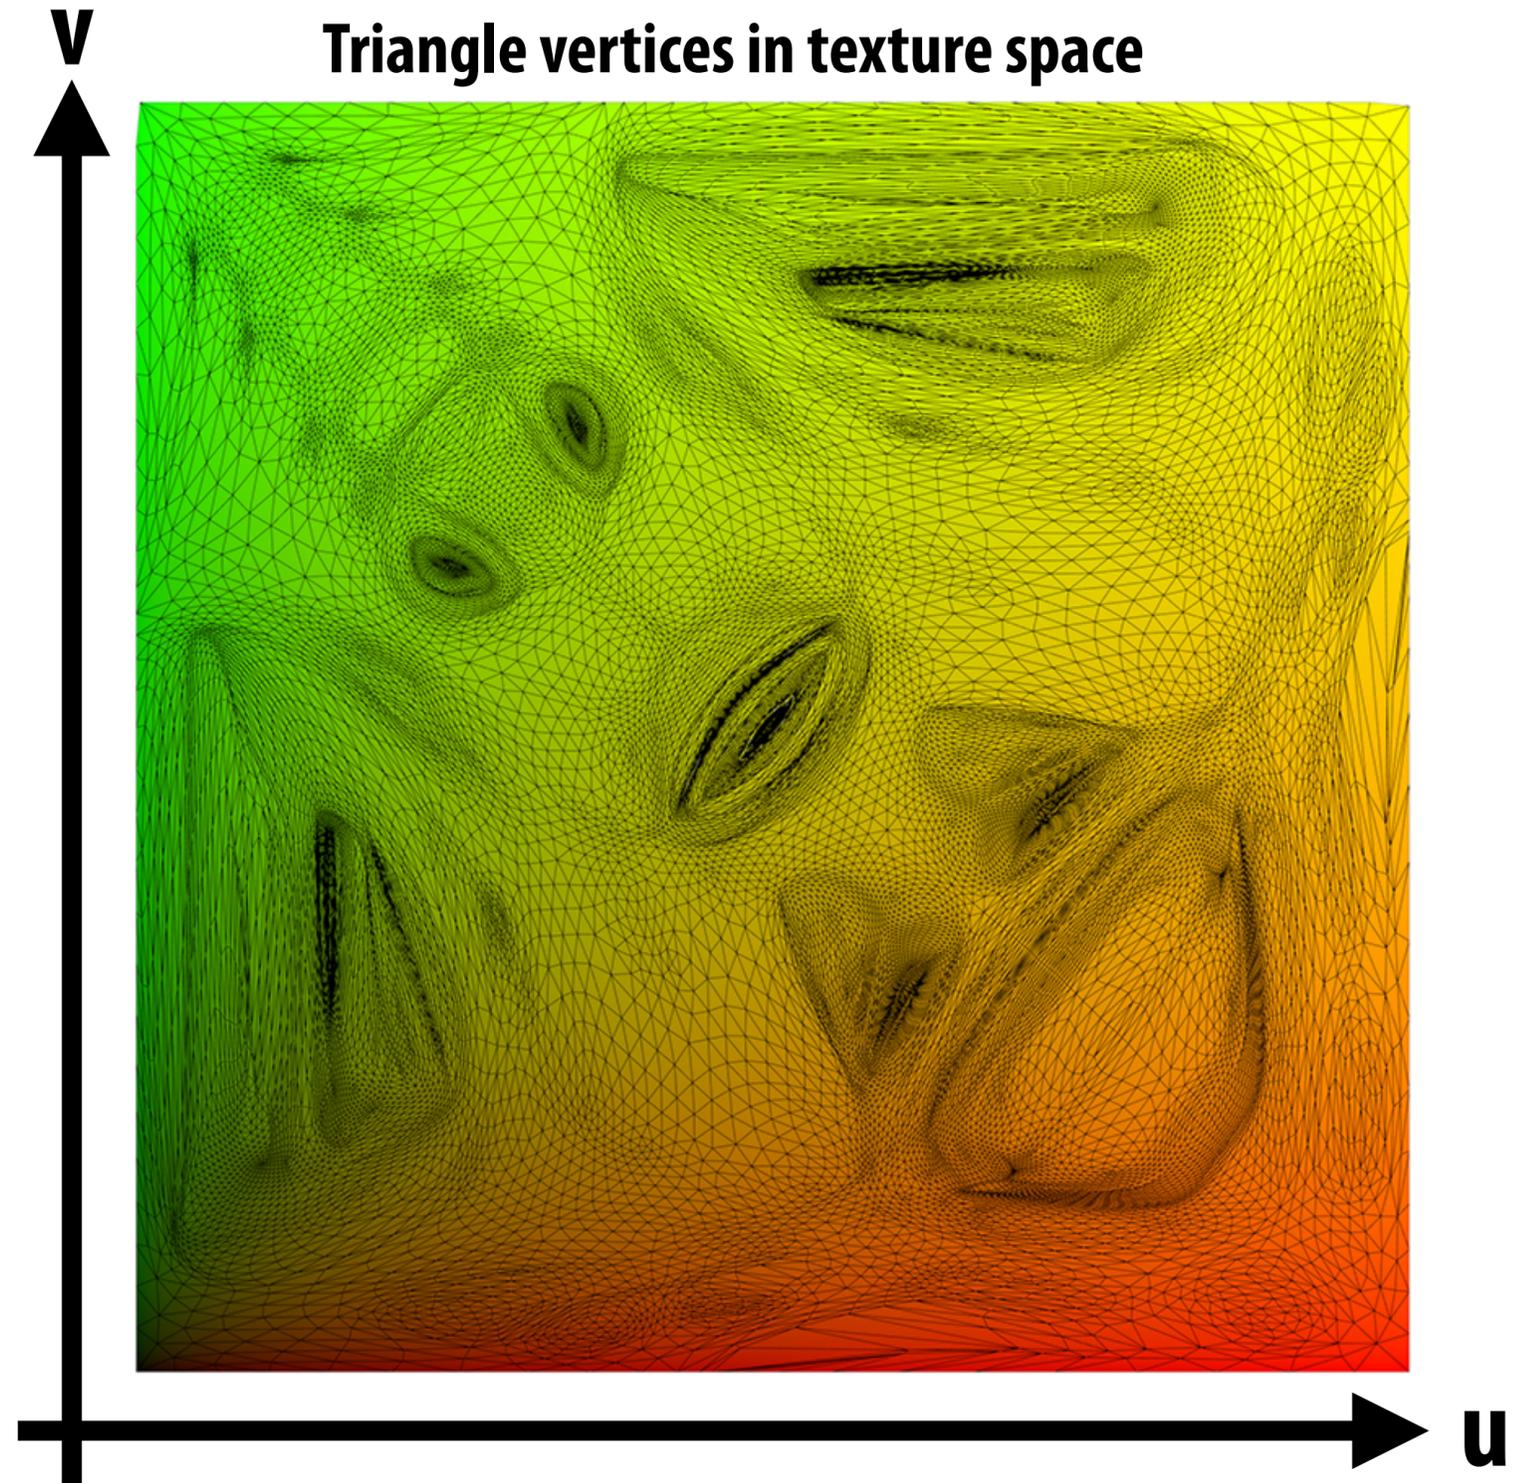


More complex mapping

Visualization of texture coordinates



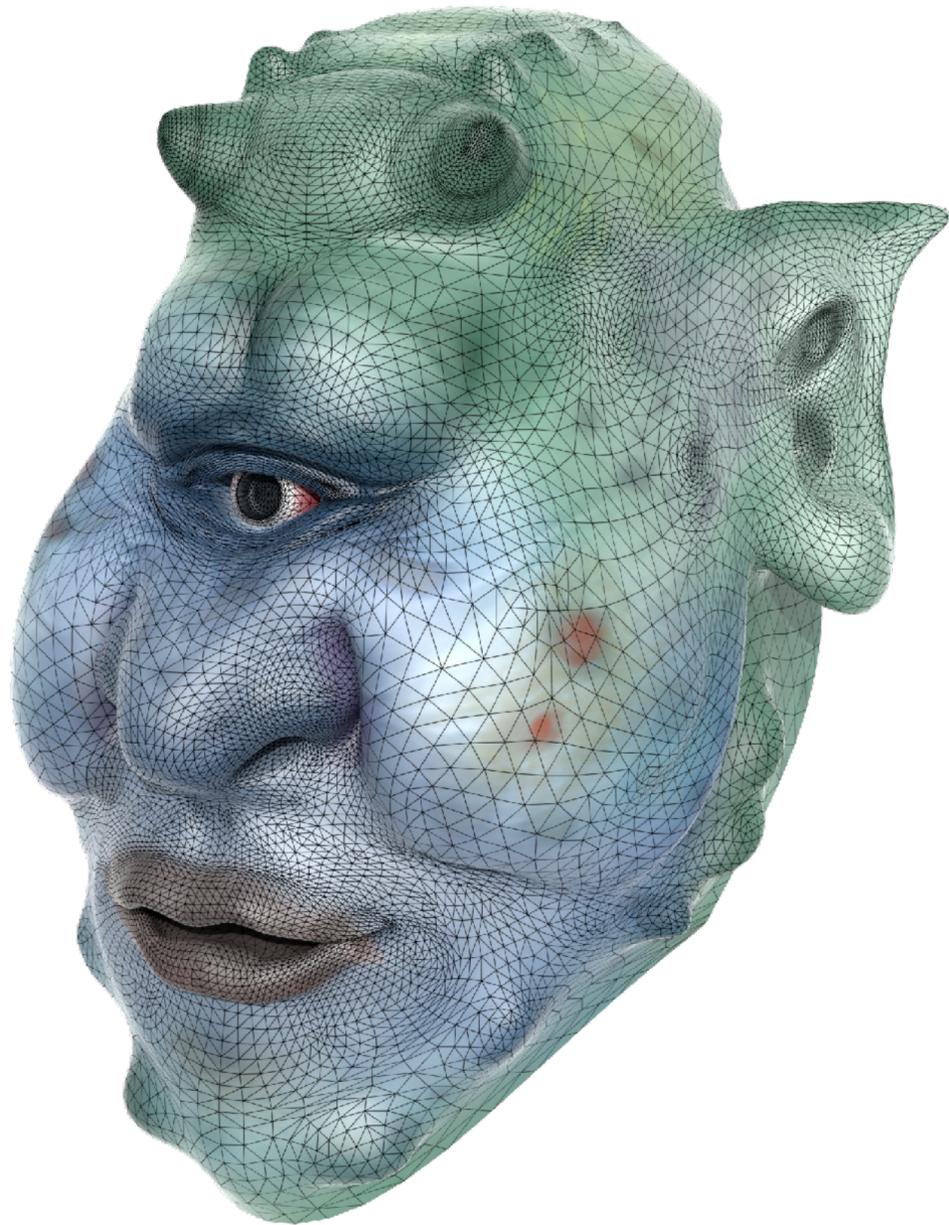
Triangle vertices in texture space



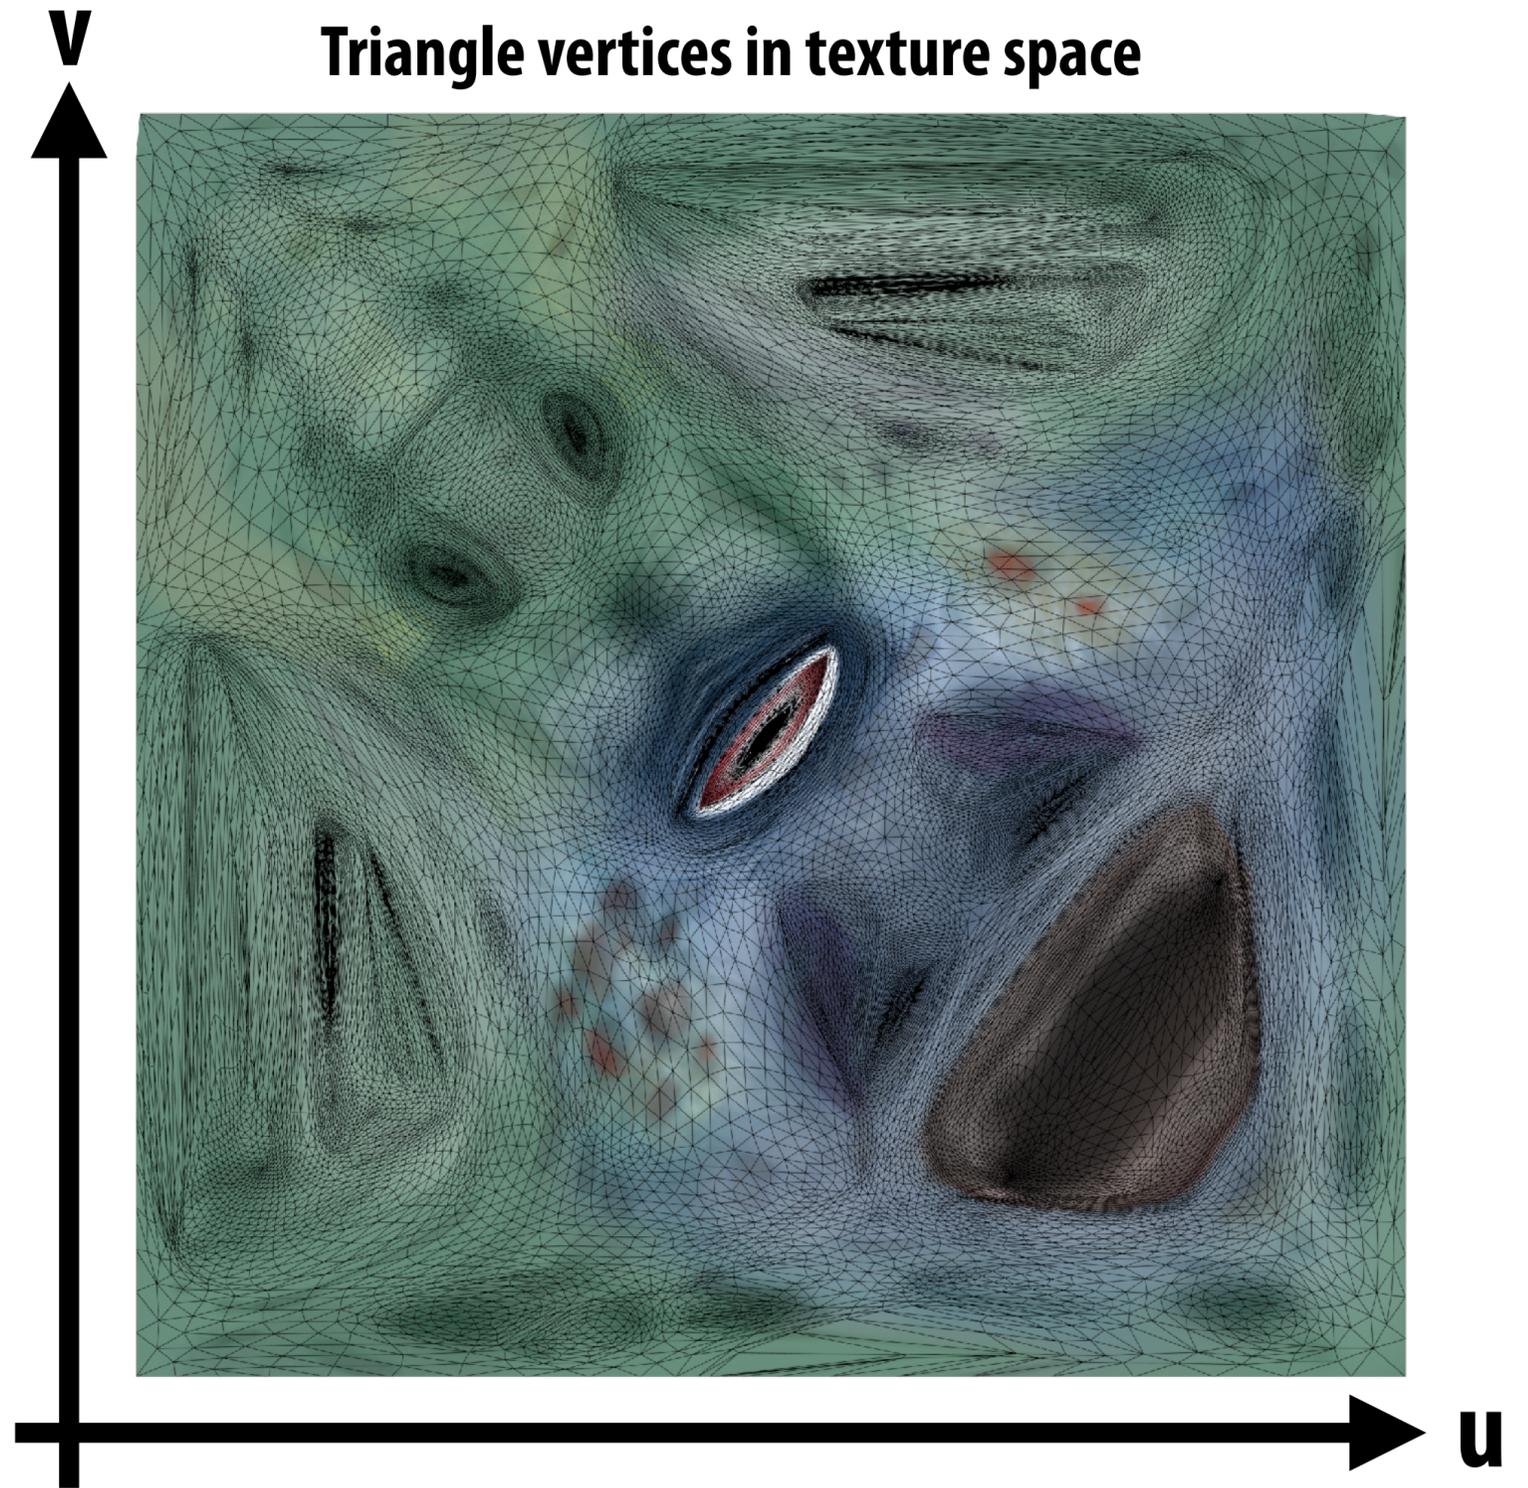
Each vertex has a coordinate (u,v) in texture space.
(Actually coming up with these coordinates is another story!)

Texture mapping adds detail

Rendered result



Triangle vertices in texture space



Texture mapping adds detail

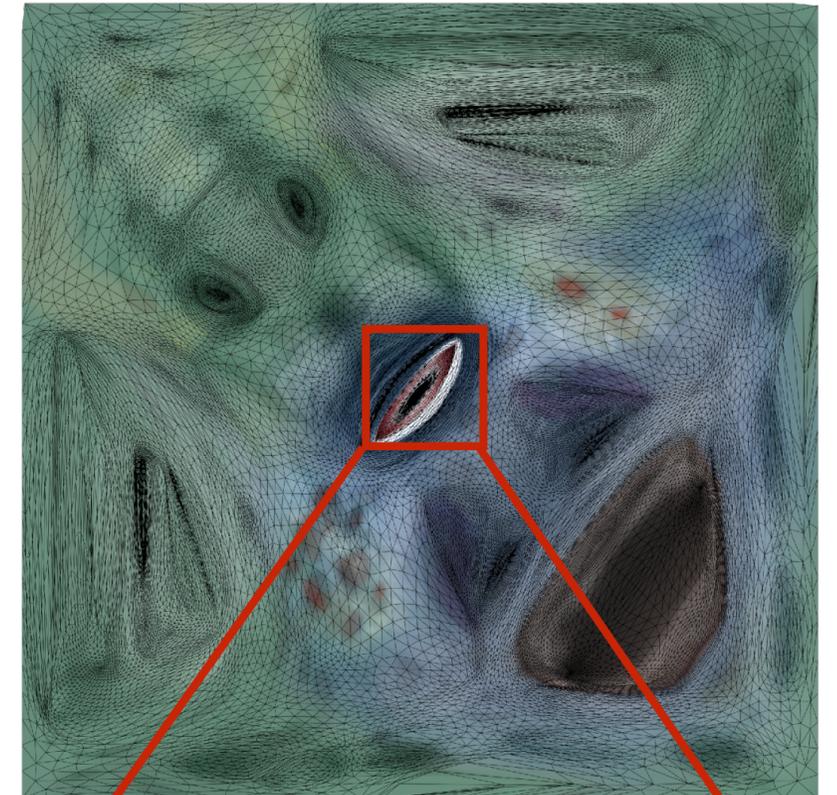
rendering without texture



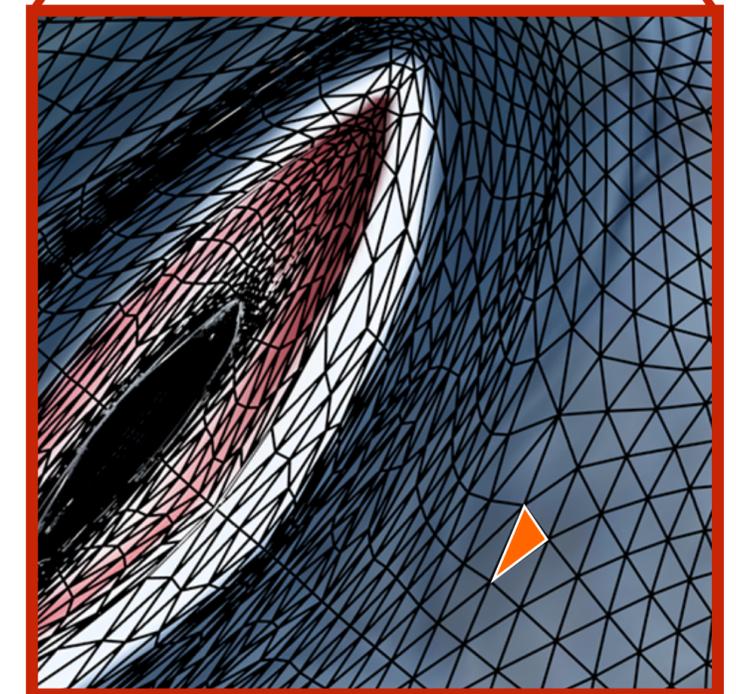
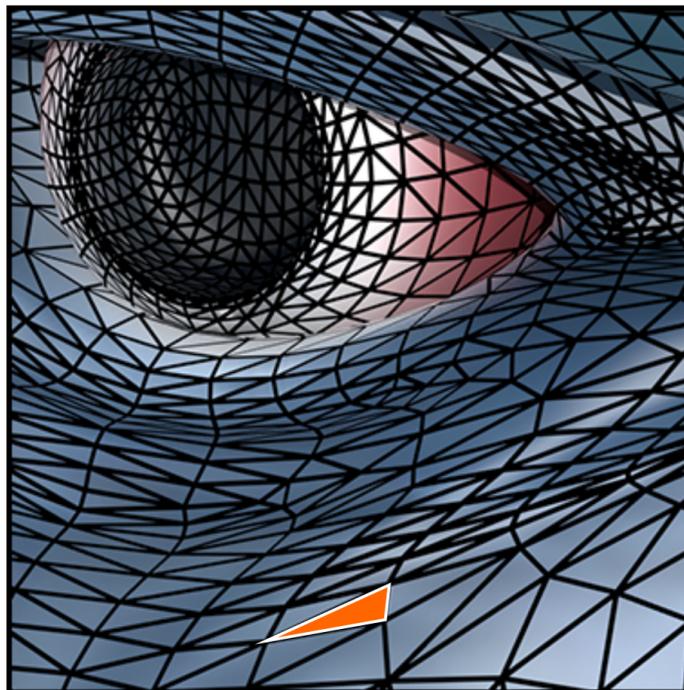
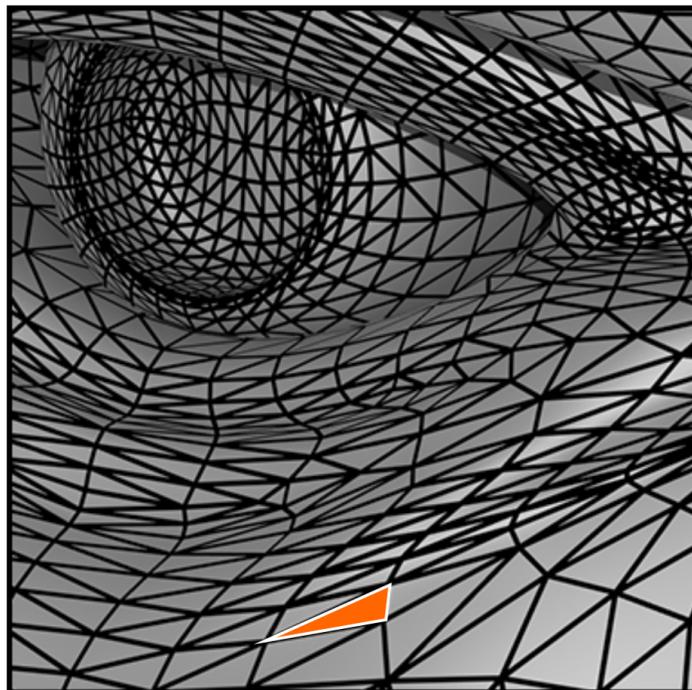
rendering with texture



texture image

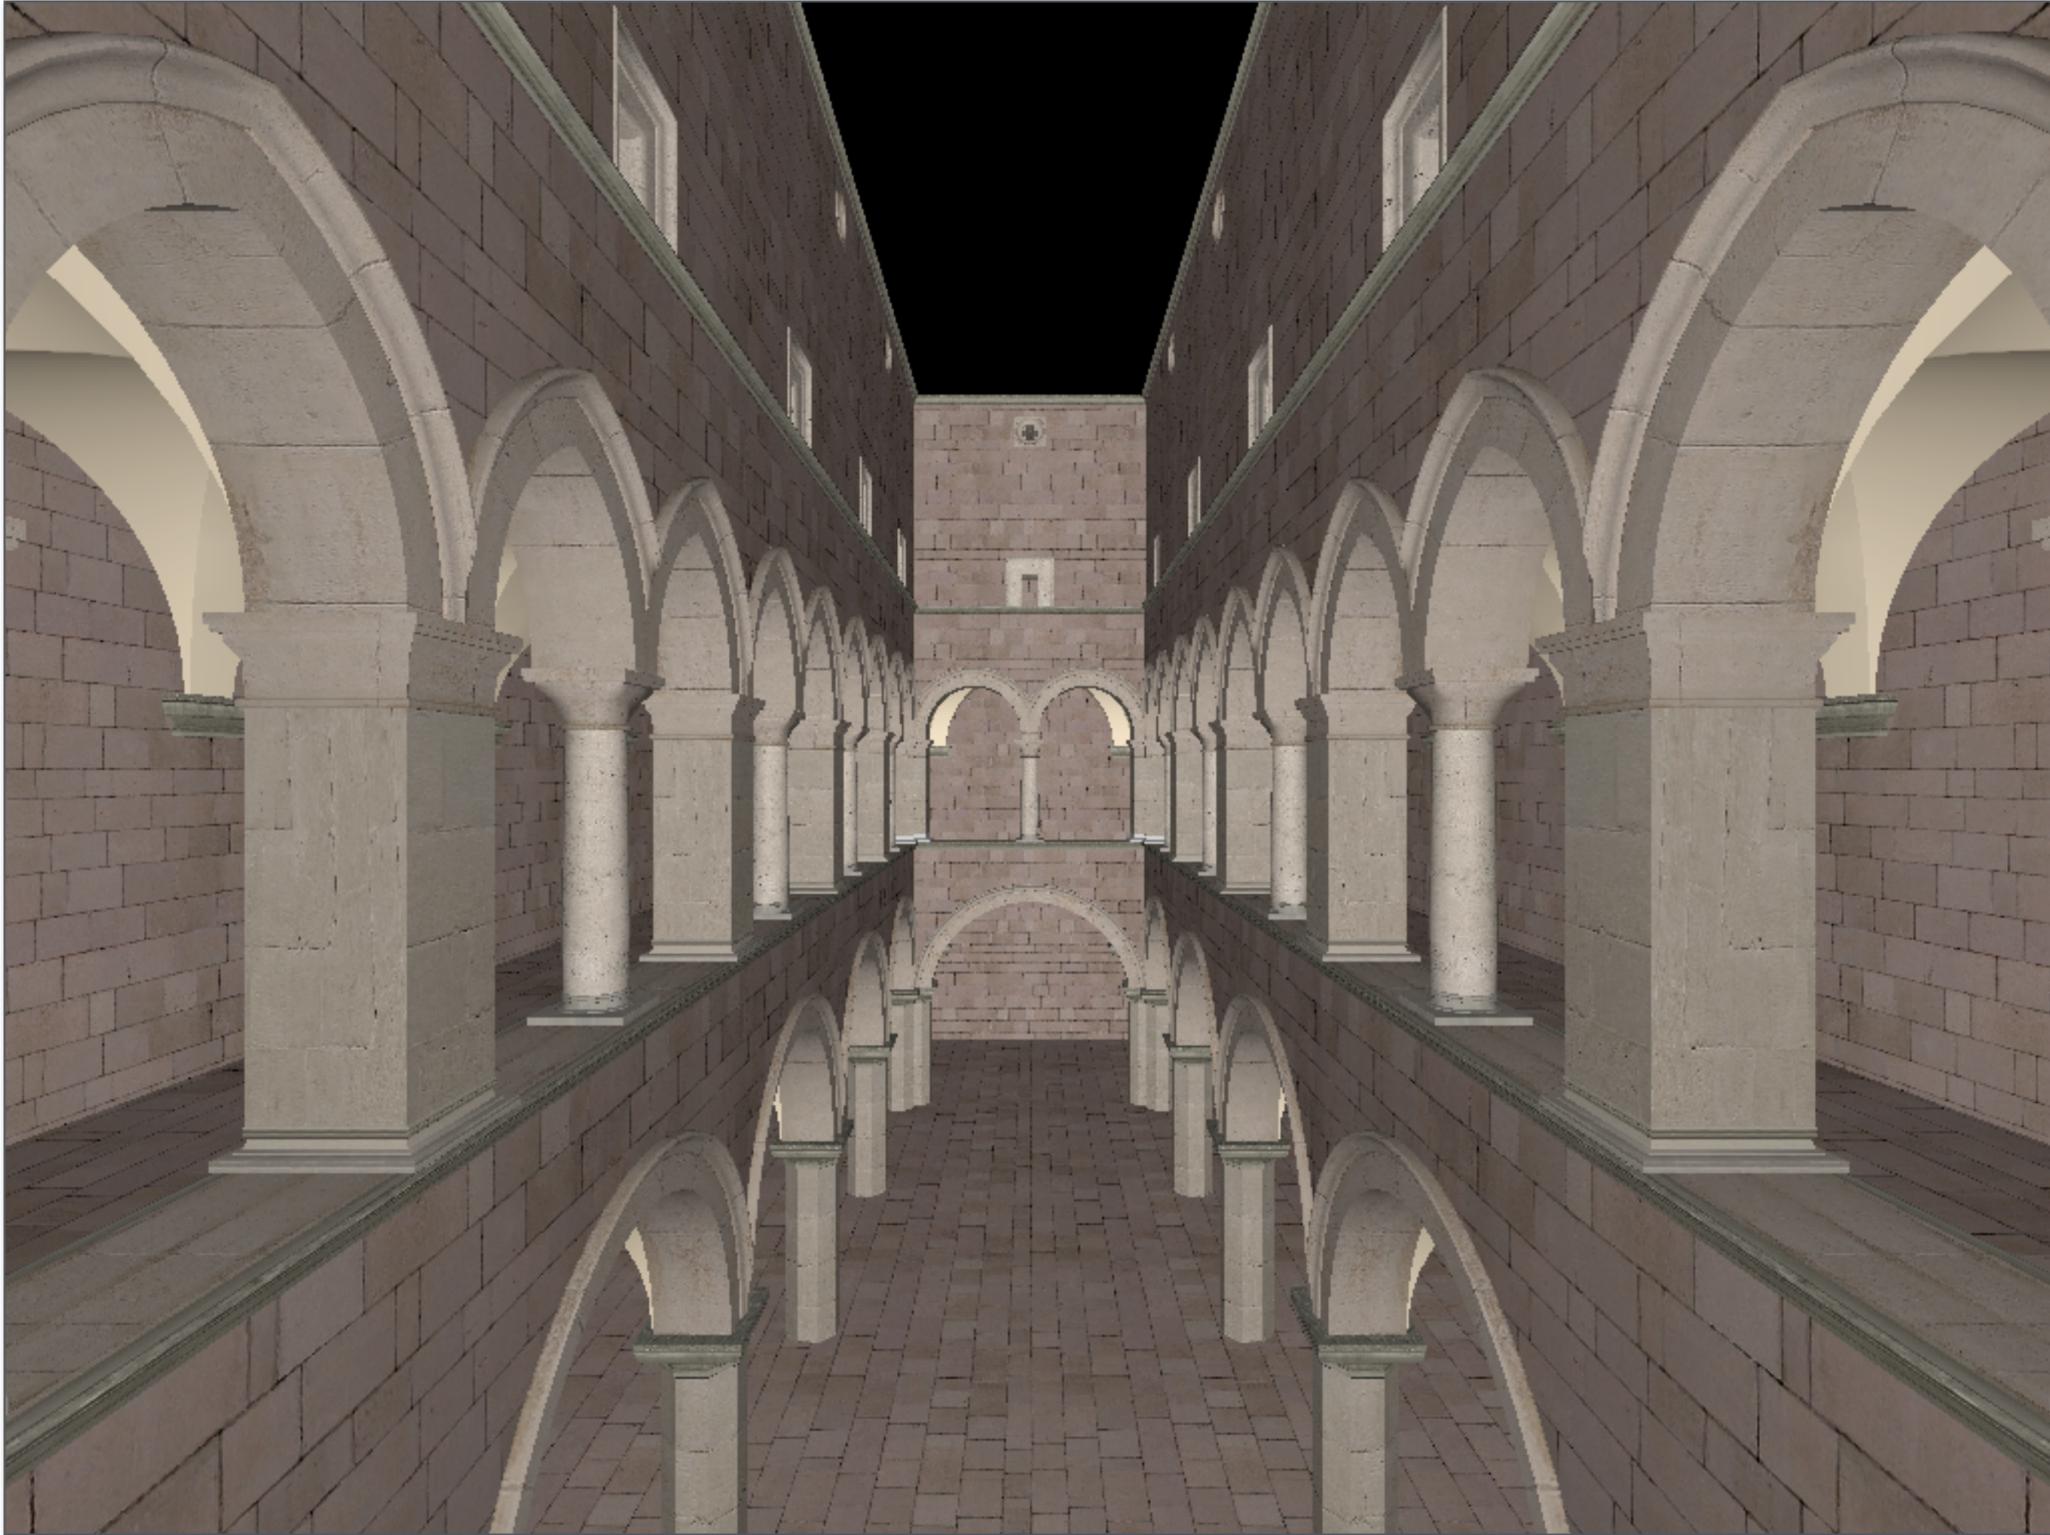


zoom

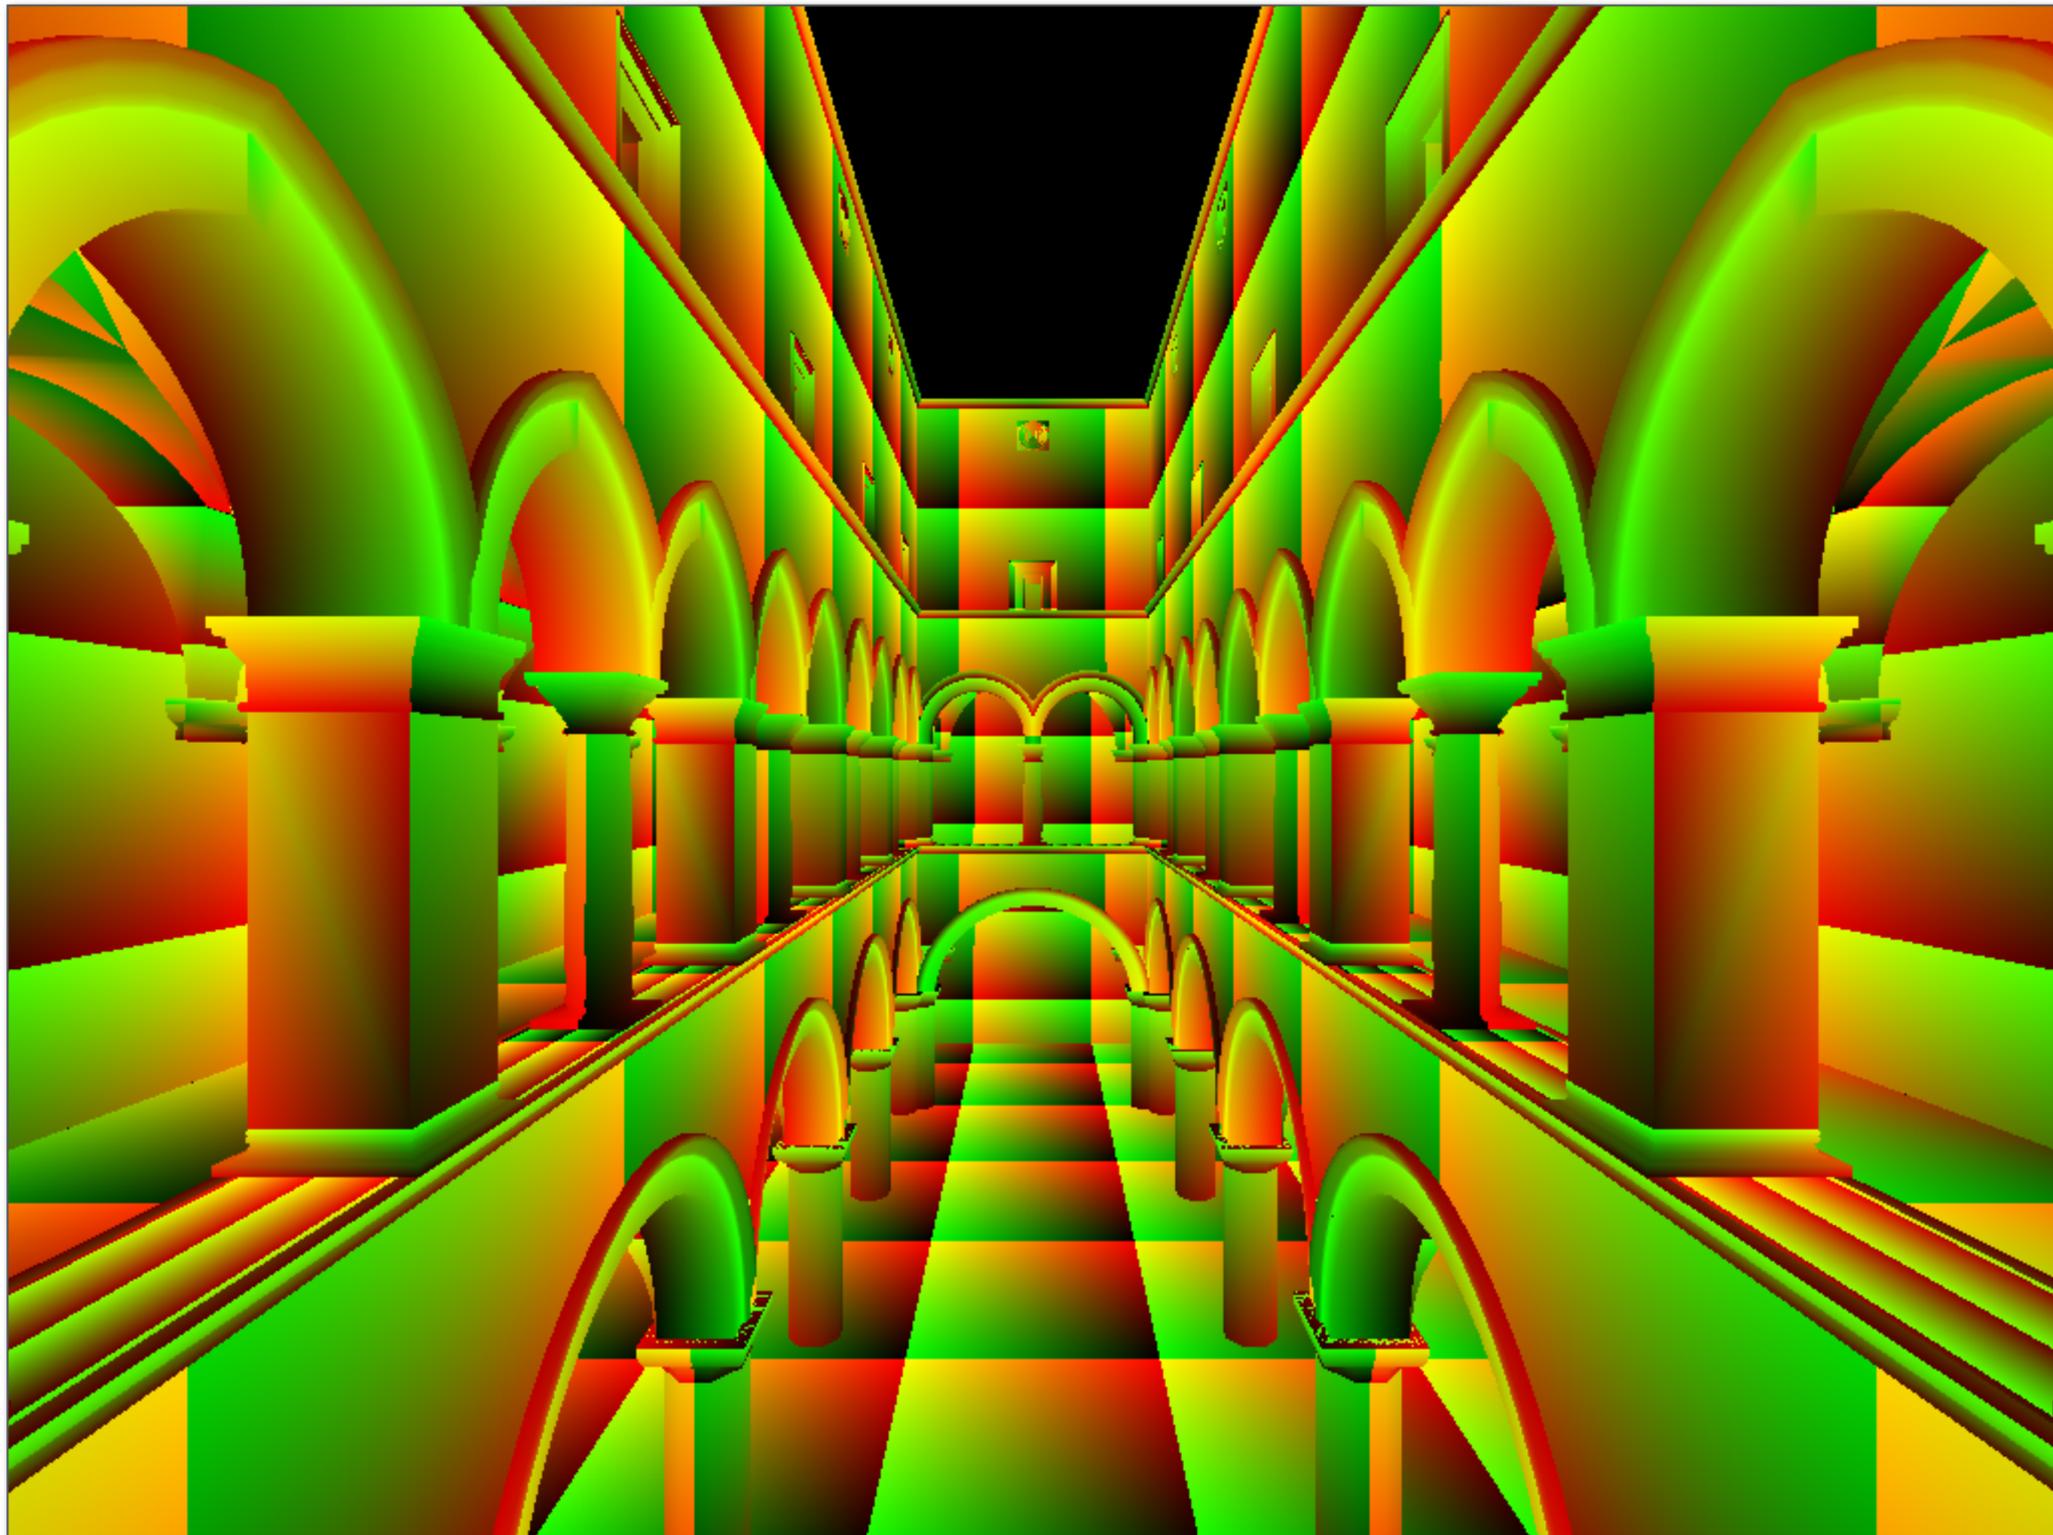


Each triangle "copies" a piece of the image back to the surface.

Textured Sponza

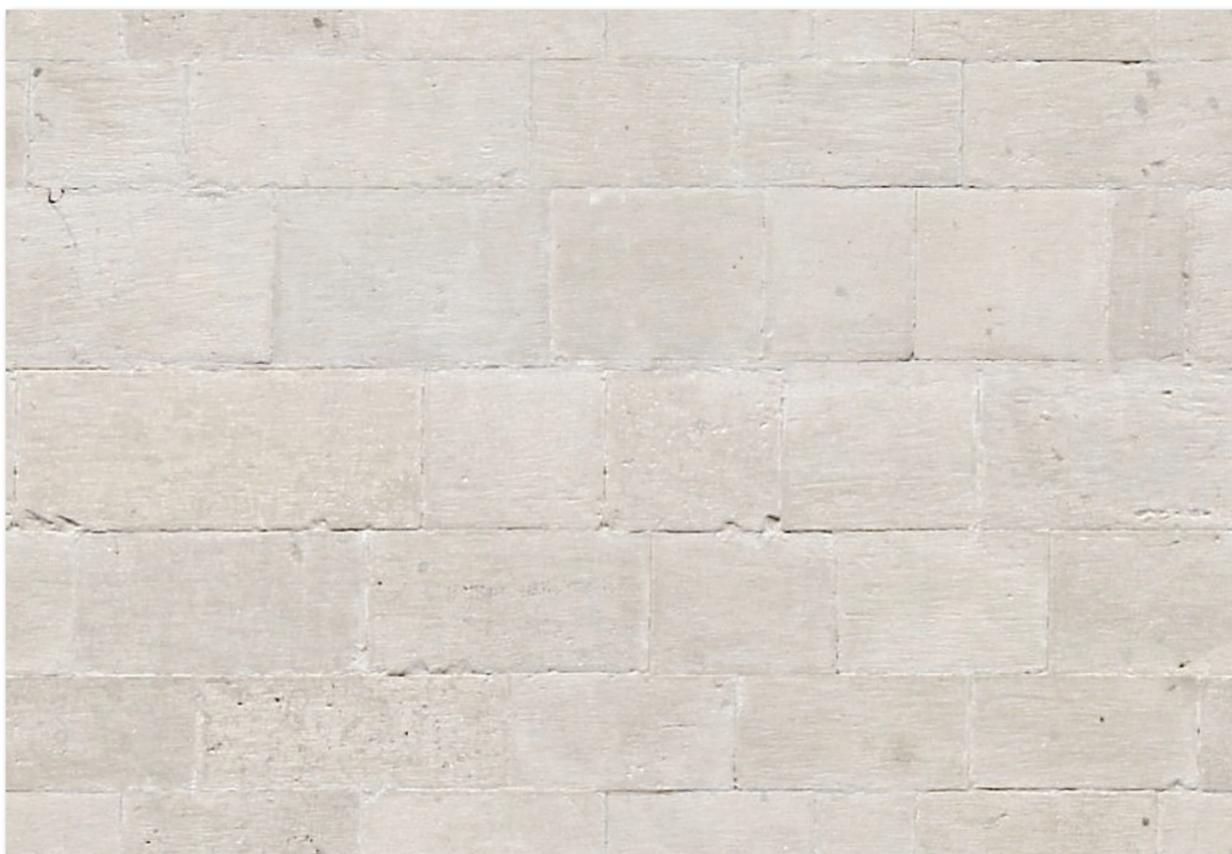


Another example: Sponza



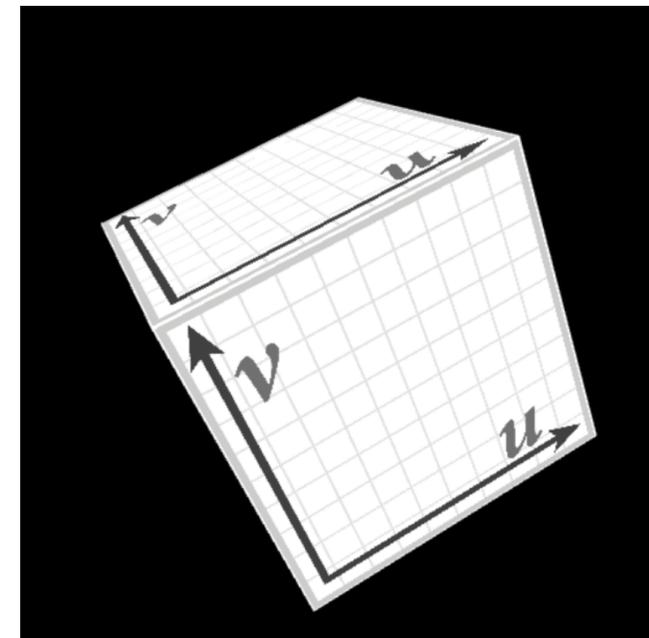
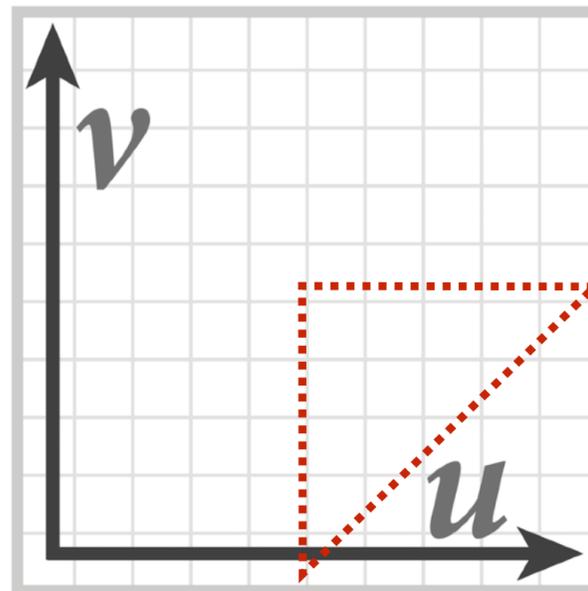
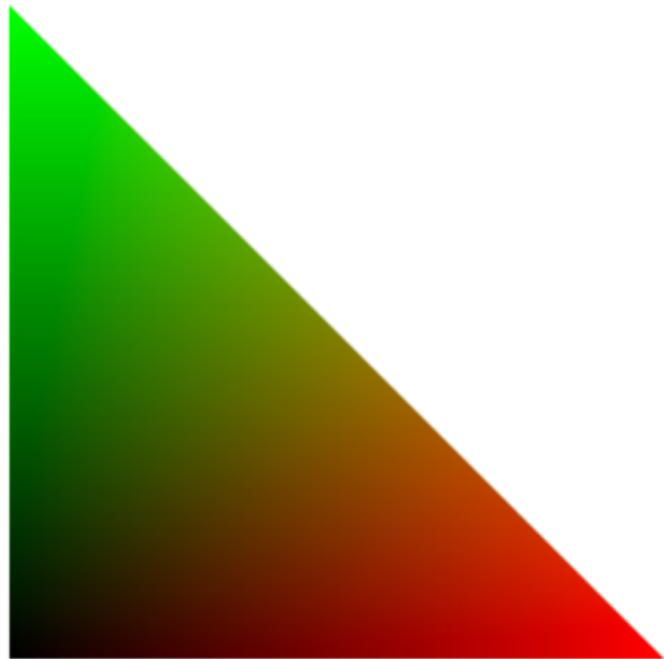
Notice texture coordinates repeat over surface.

Example textures used in Sponza



Texture sampling 101

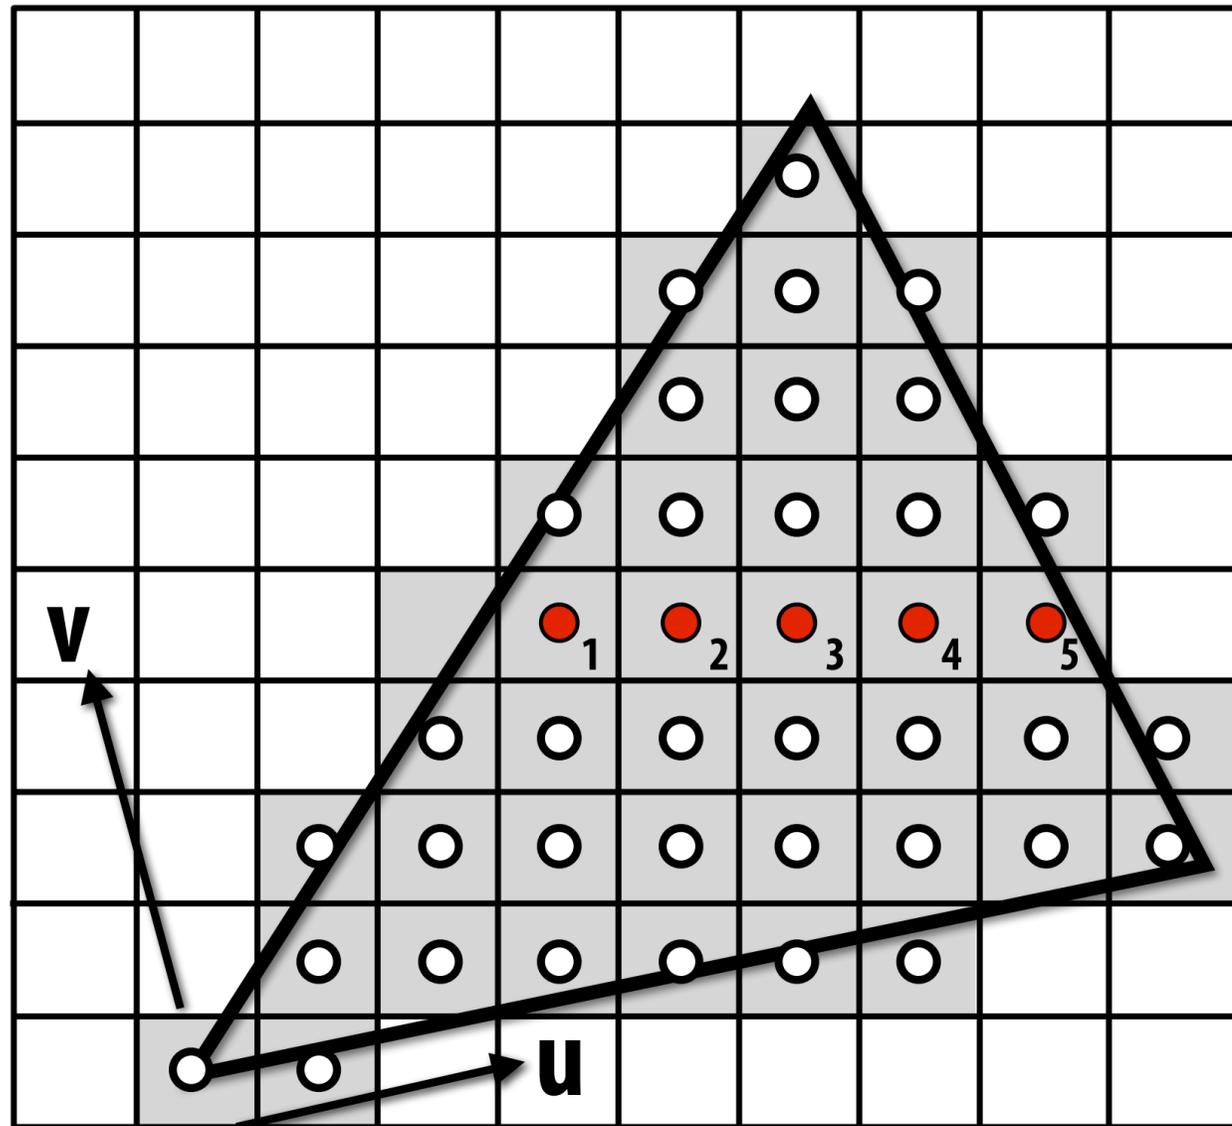
- **Basic algorithm for mapping texture to surface:**
 - **Interpolate U and V coordinates across triangle**
 - **For each fragment**
 - **Sample (evaluate) texture at (U,V)**
 - **Set color of fragment to sampled texture value**



...sadly not this easy in general!

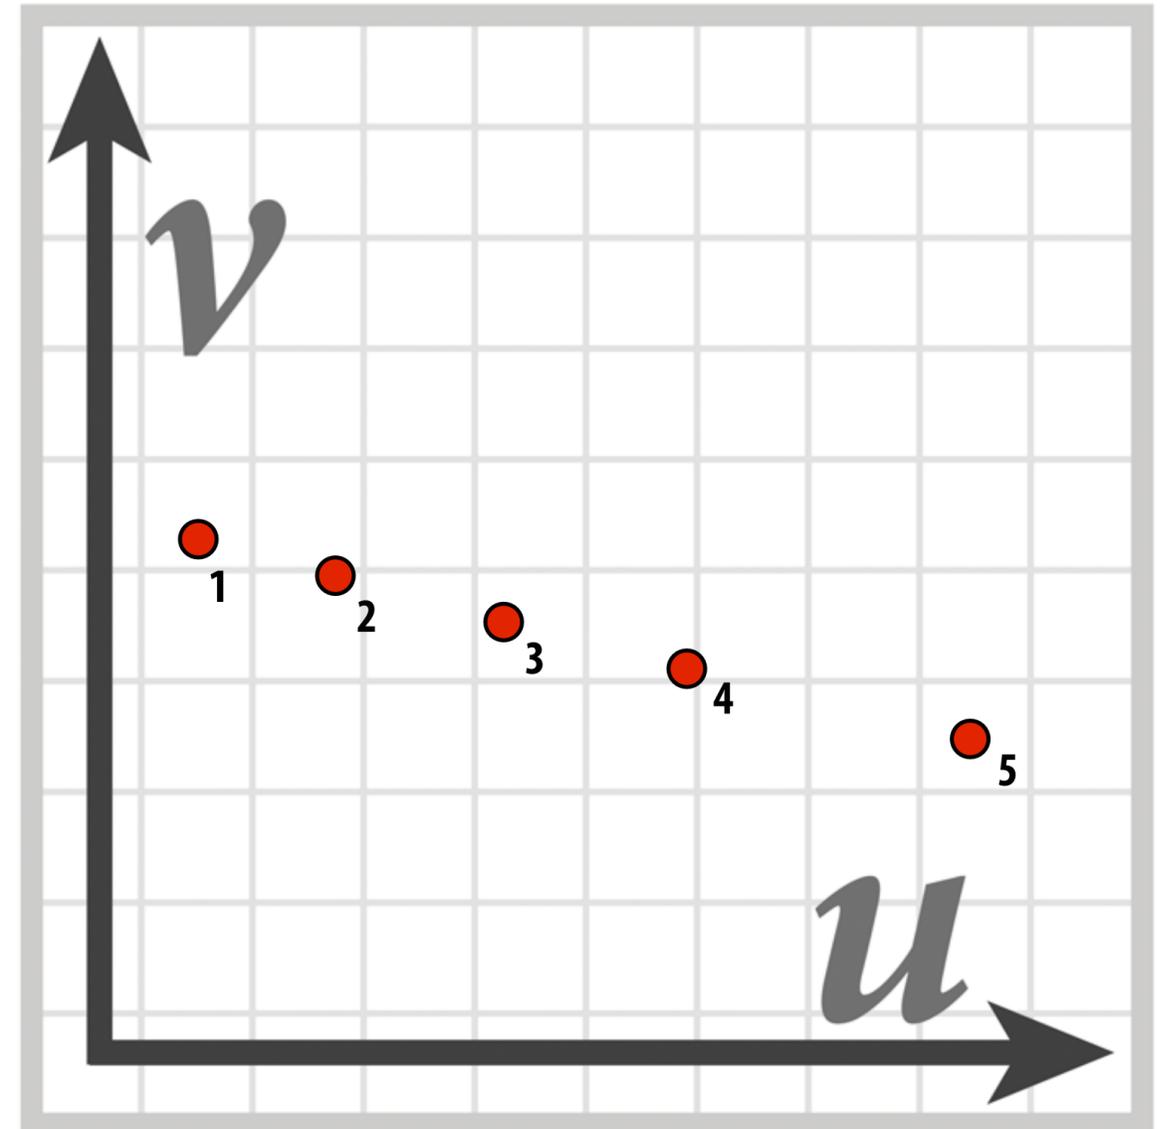
Texture space samples

Sample positions in XY screen space



Sample positions are uniformly distributed in screen space (rasterizer samples triangle's appearance at these locations)

Sample positions in texture space



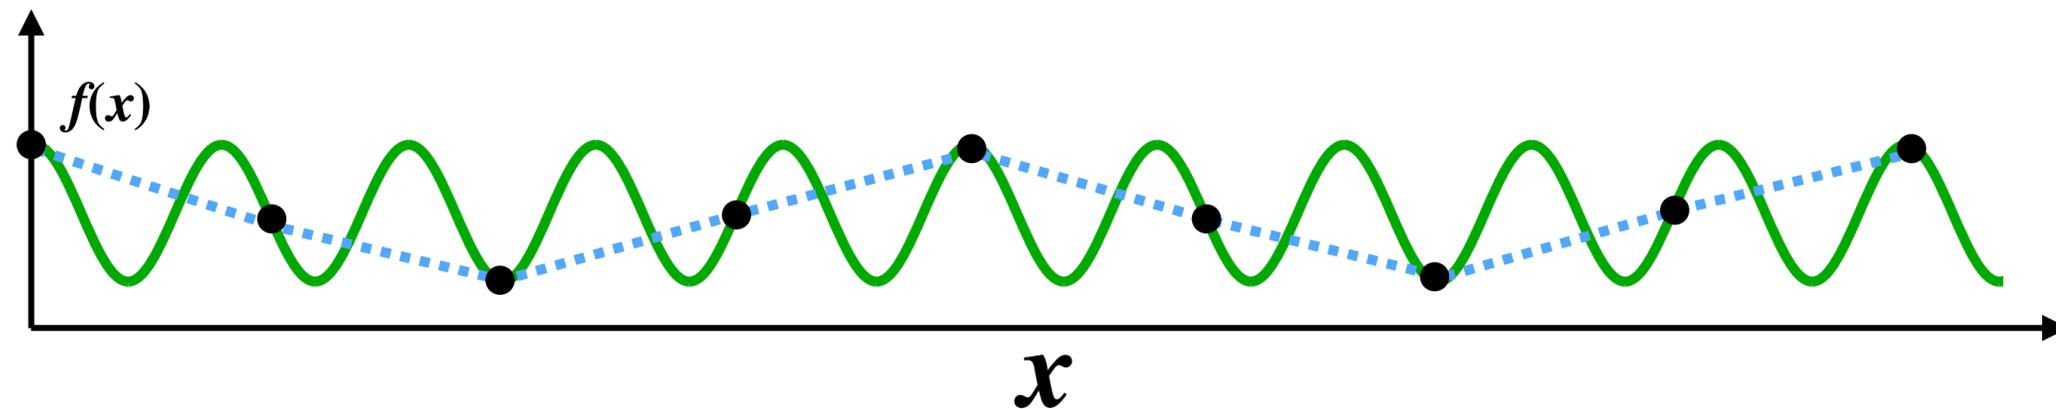
Texture sample positions in texture space (texture function is sampled at these locations)

Applying textures is a form of sampling!

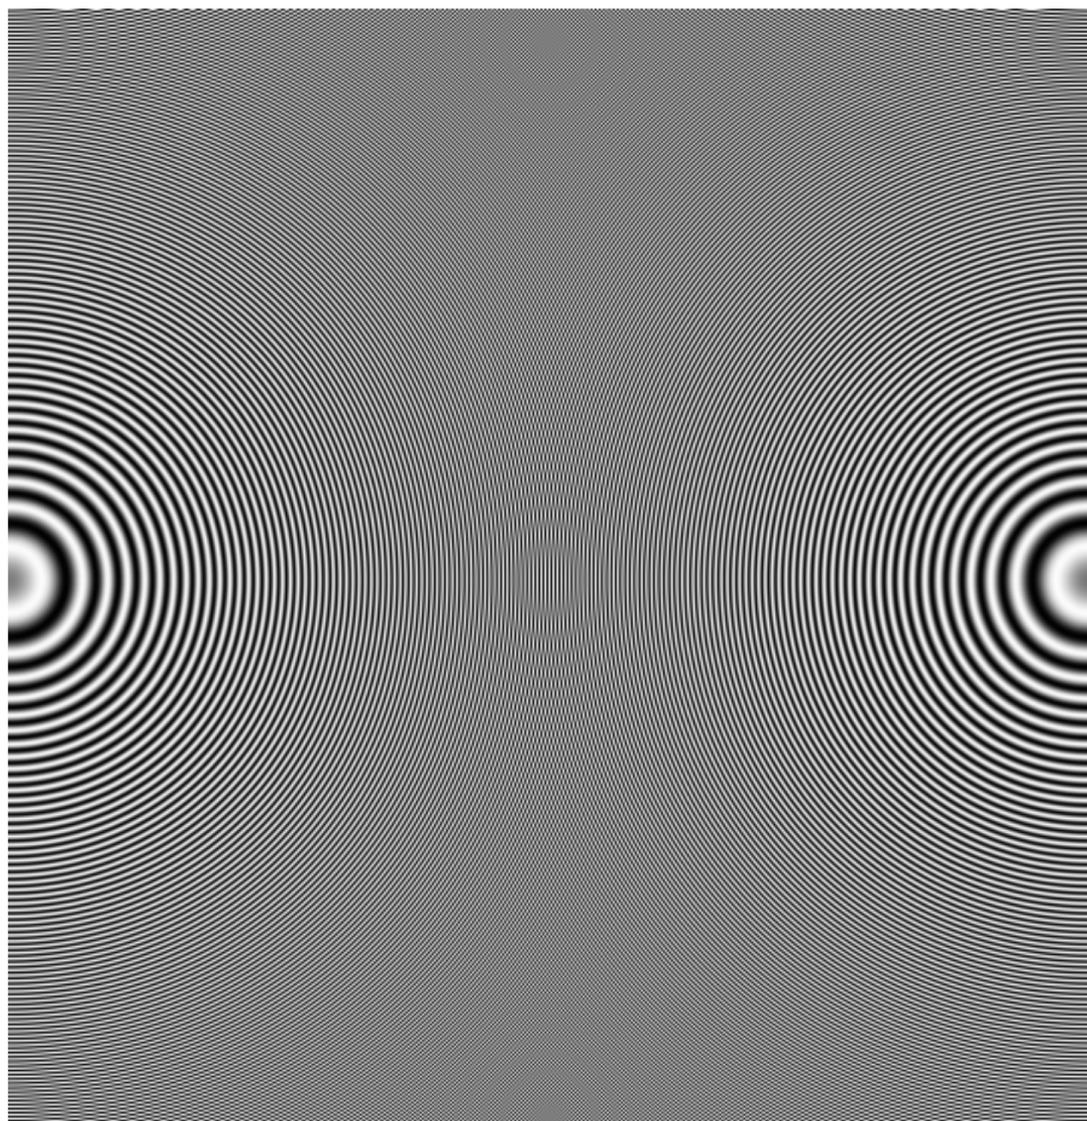
$t(u,v)$

Recall: aliasing

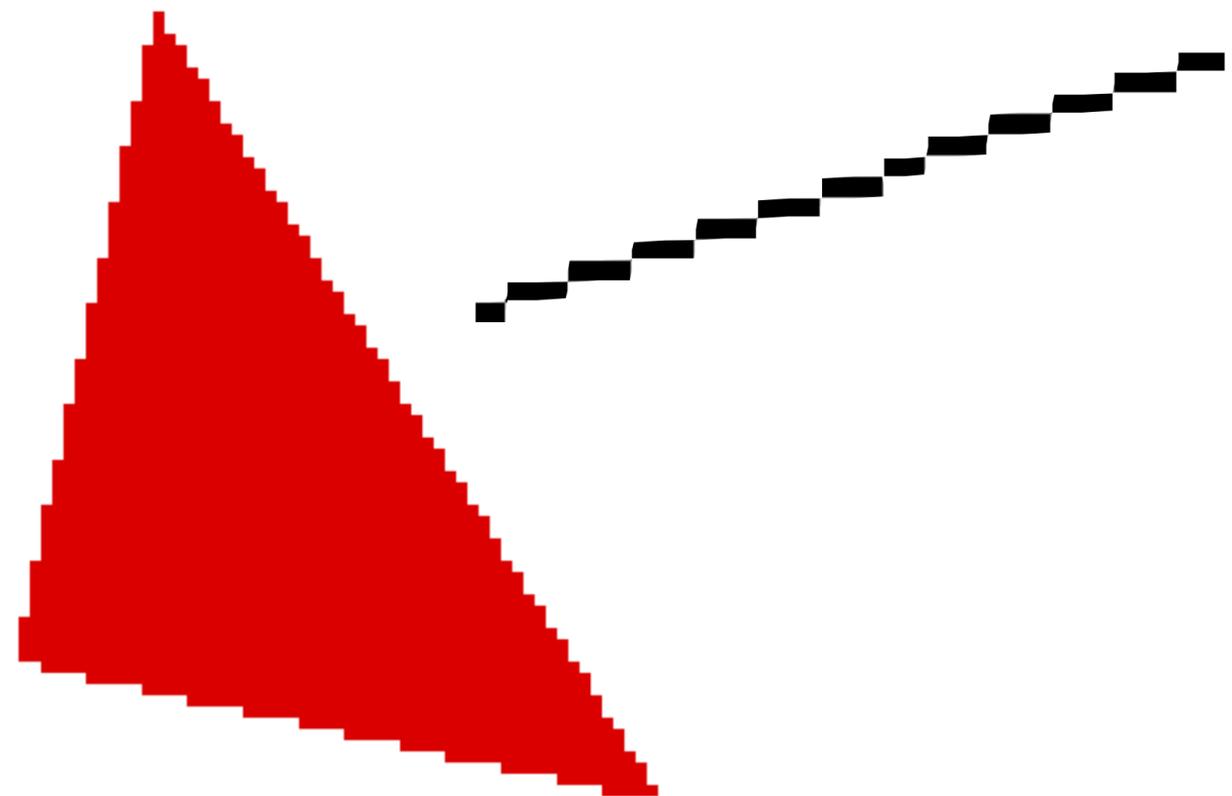
Undersampling a high-frequency signal can result in aliasing



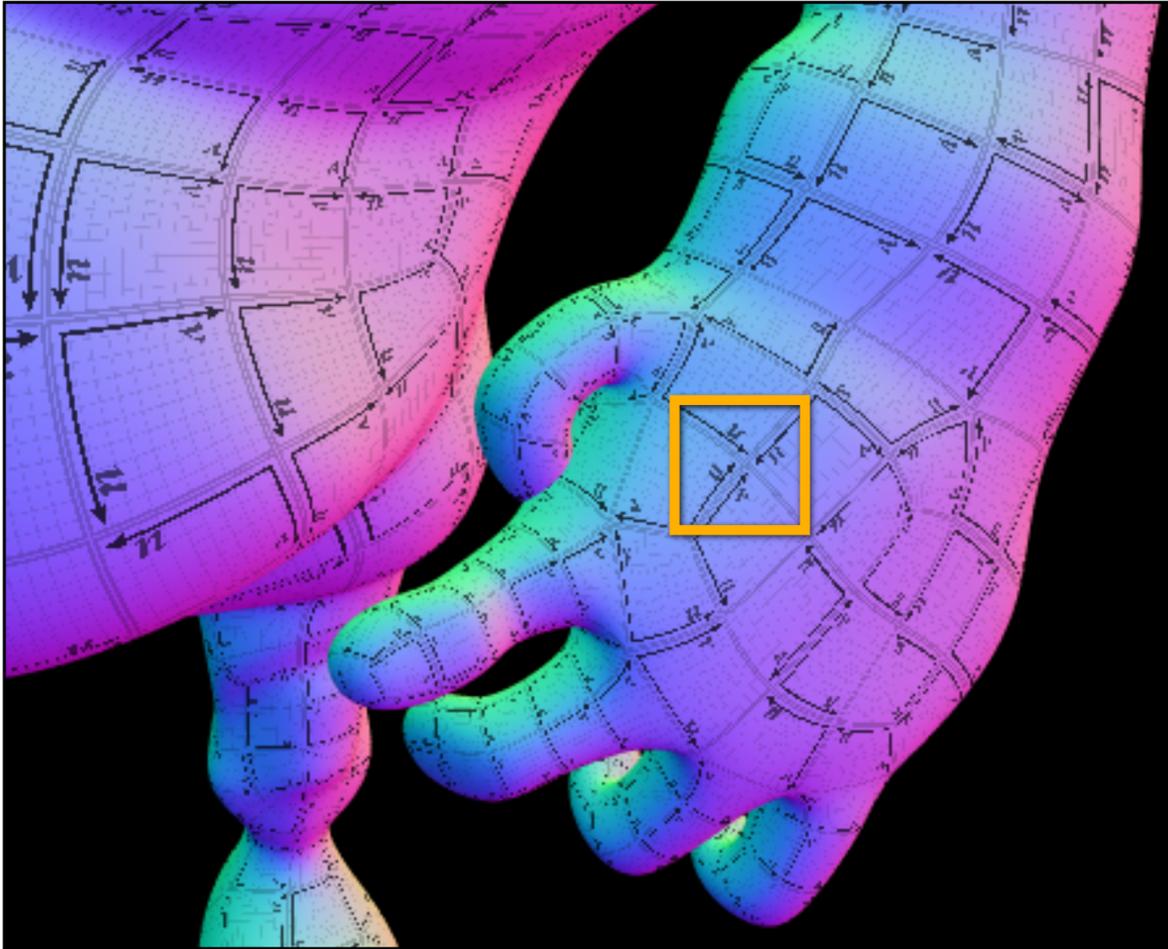
1D example



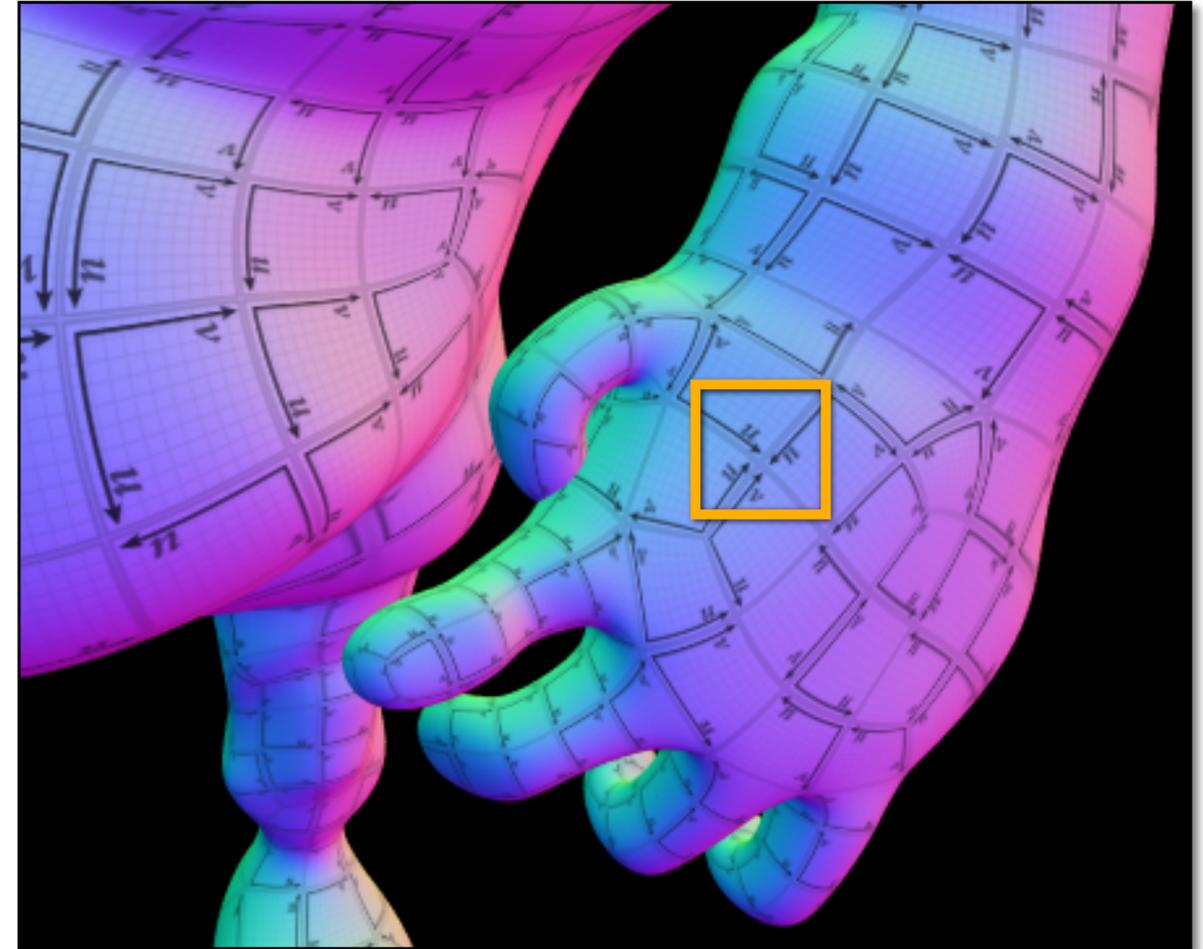
2D examples:
Moiré patterns, jaggies



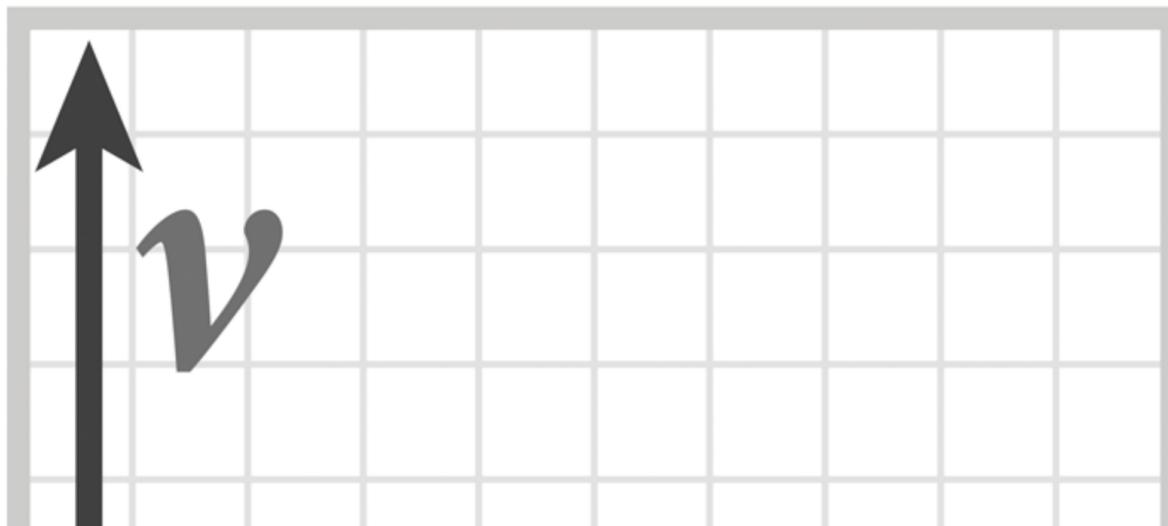
Aliasing due to undersampling texture



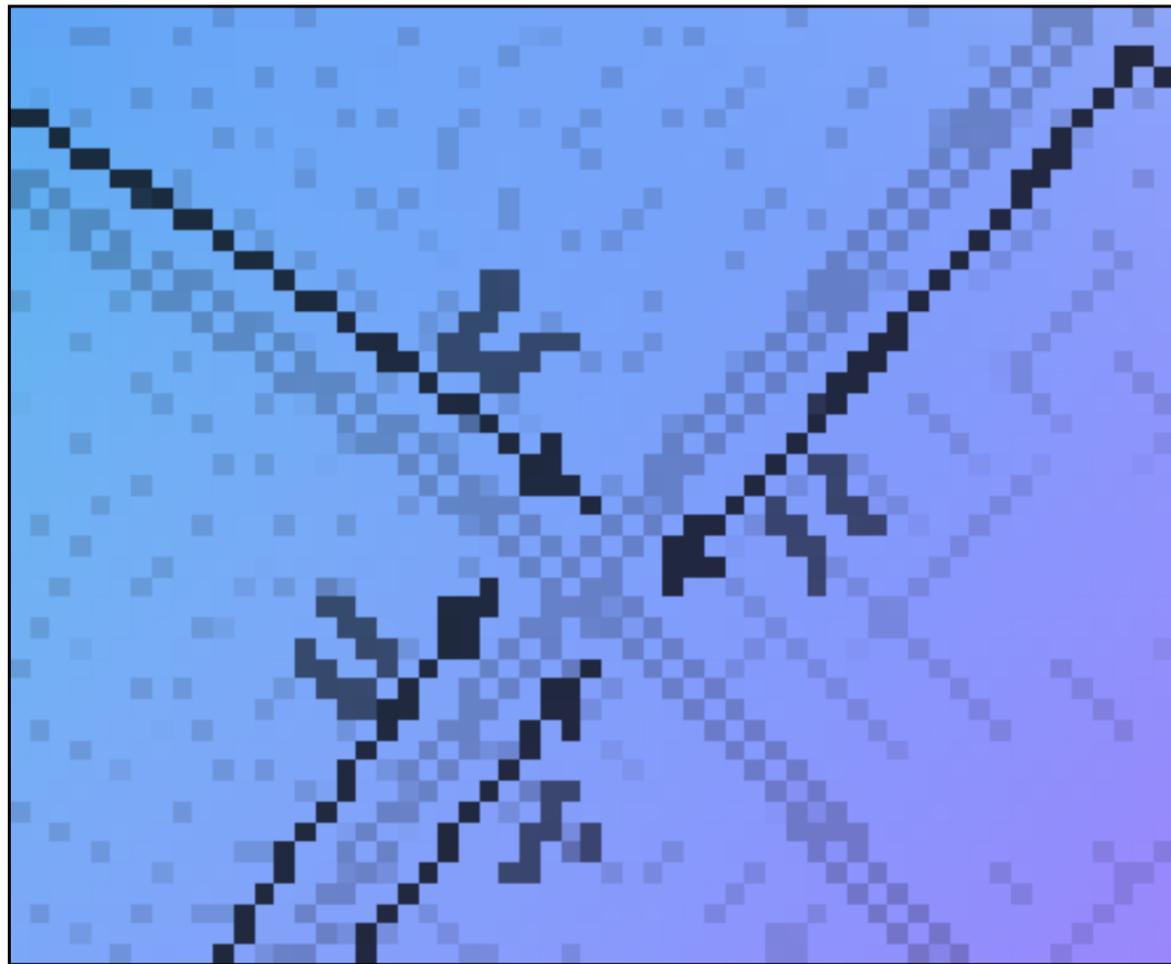
No pre-filtering of texture data
(resulting image exhibits aliasing)



Rendering using pre-filtered texture data



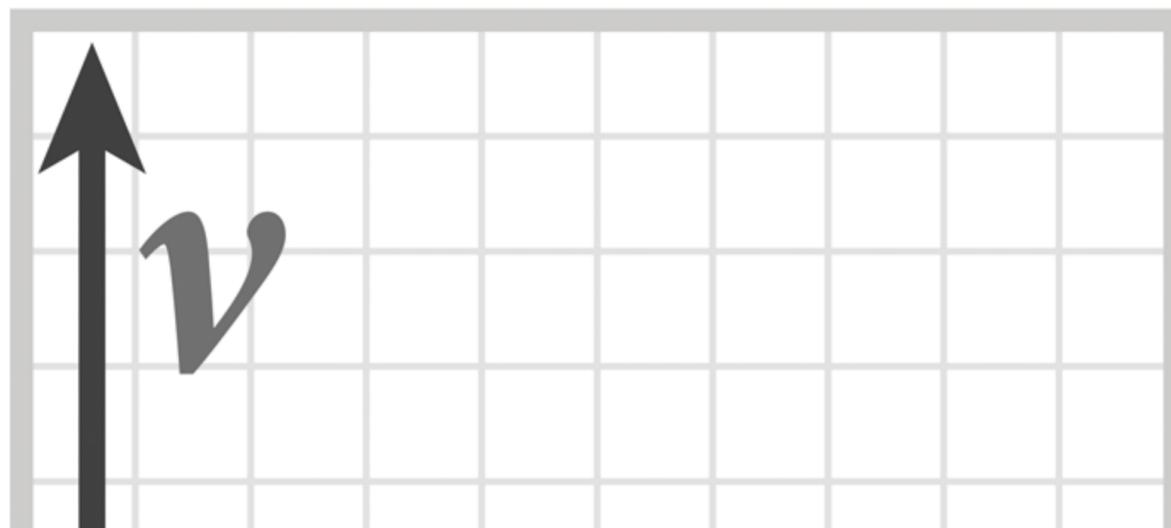
Aliasing due to undersampling (zoom)



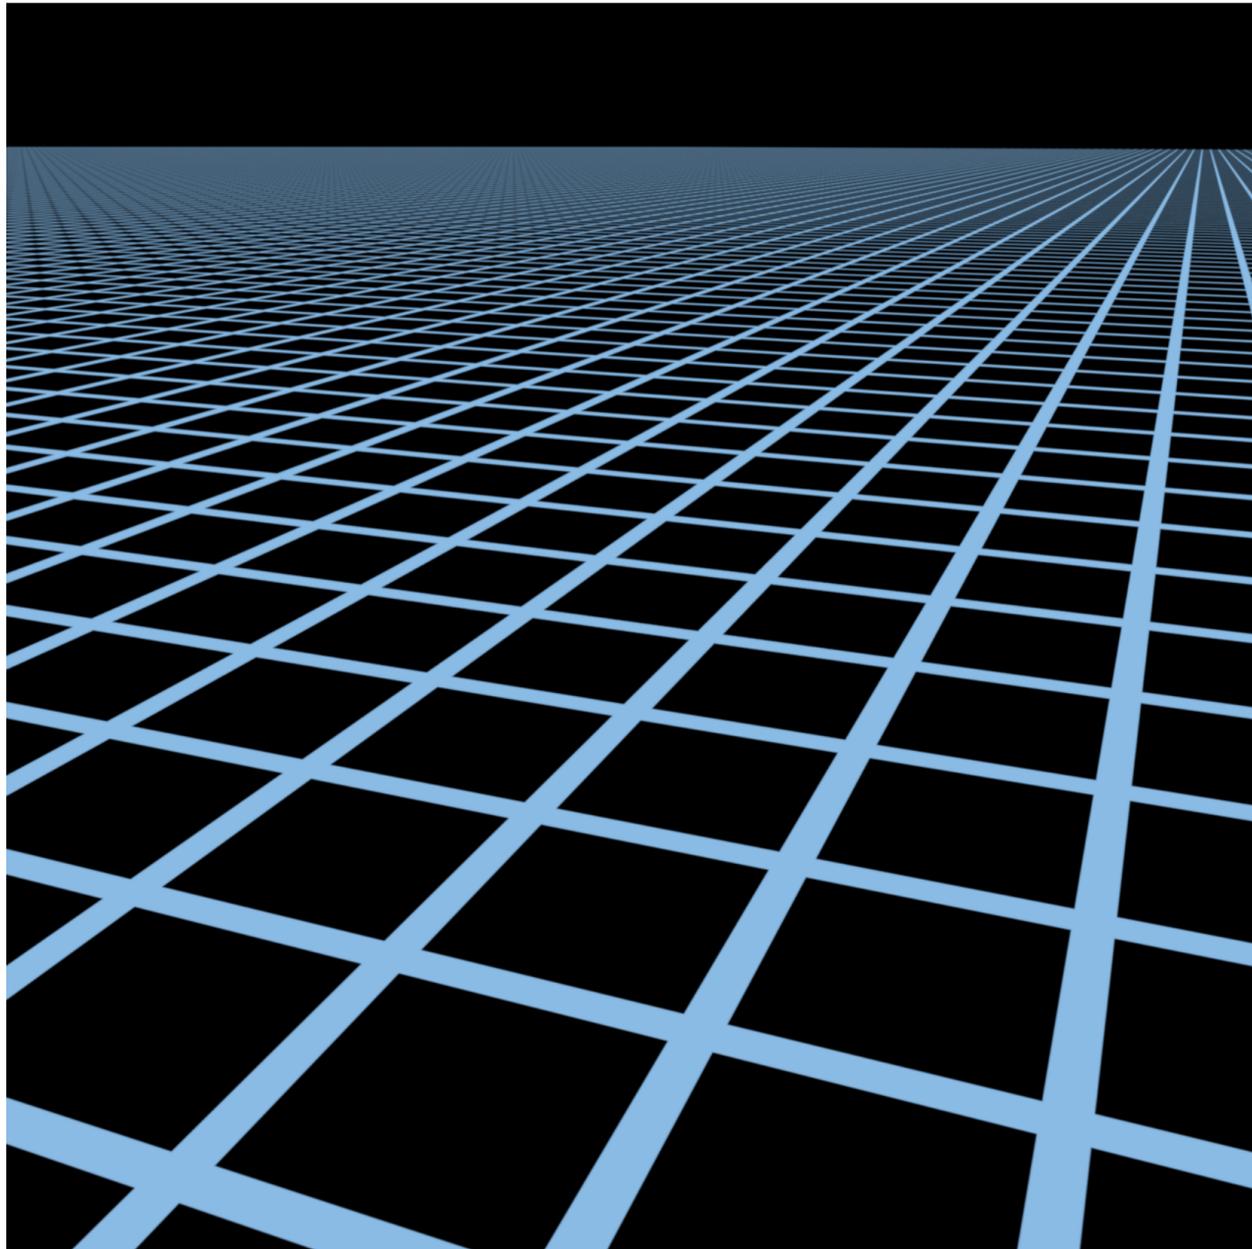
No pre-filtering of texture data
(resulting image exhibits aliasing)



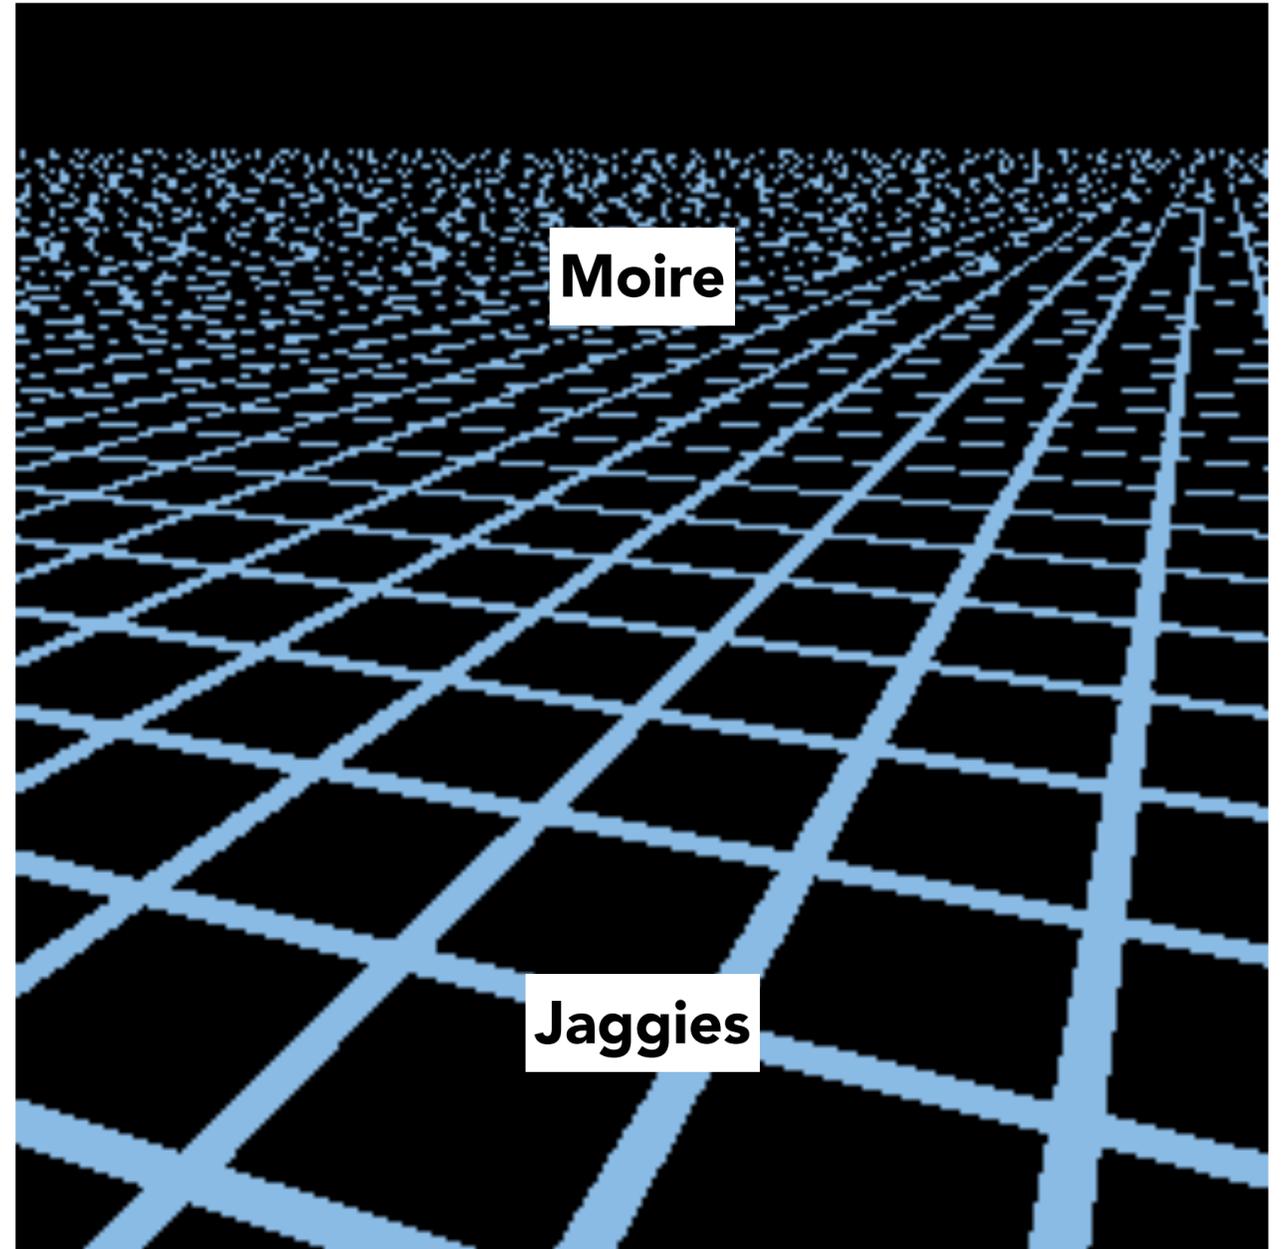
Rendering using pre-filtered texture data



Point sampling textures



Source image: 1280x1280 pixels

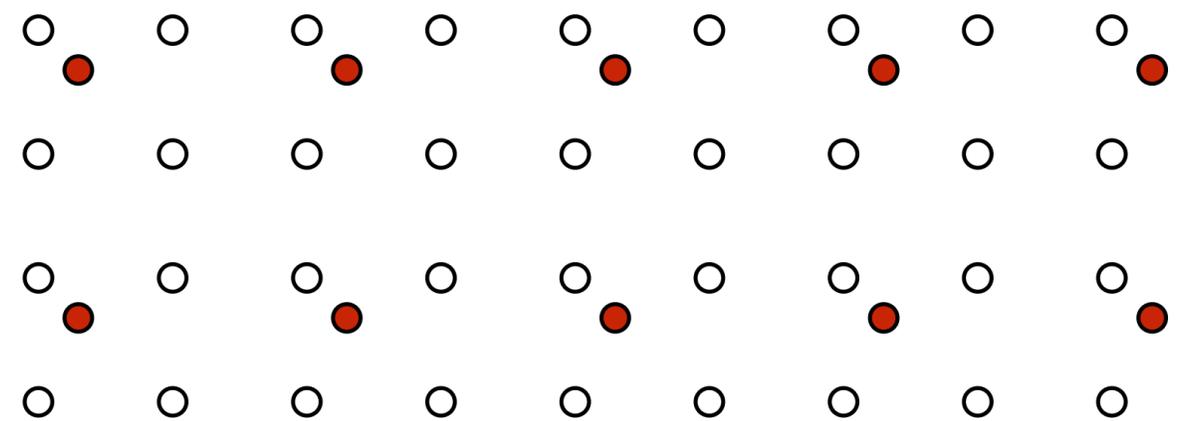
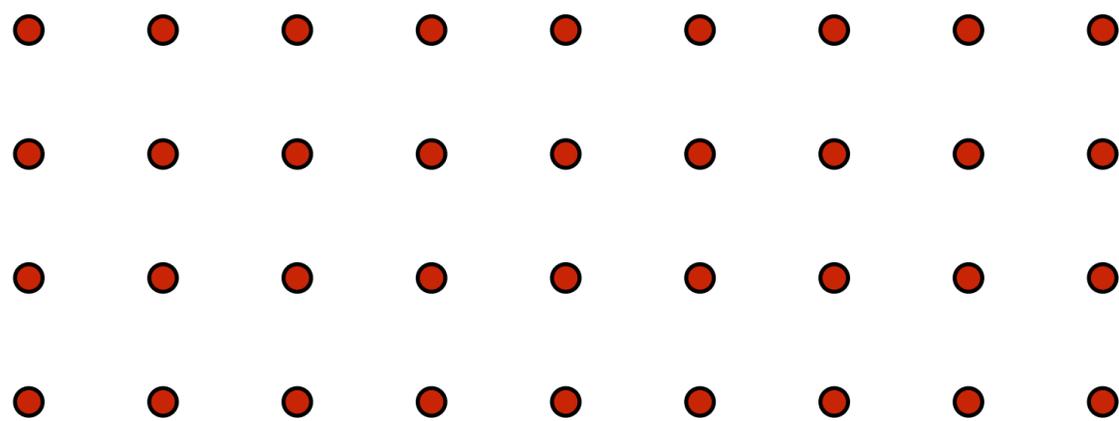
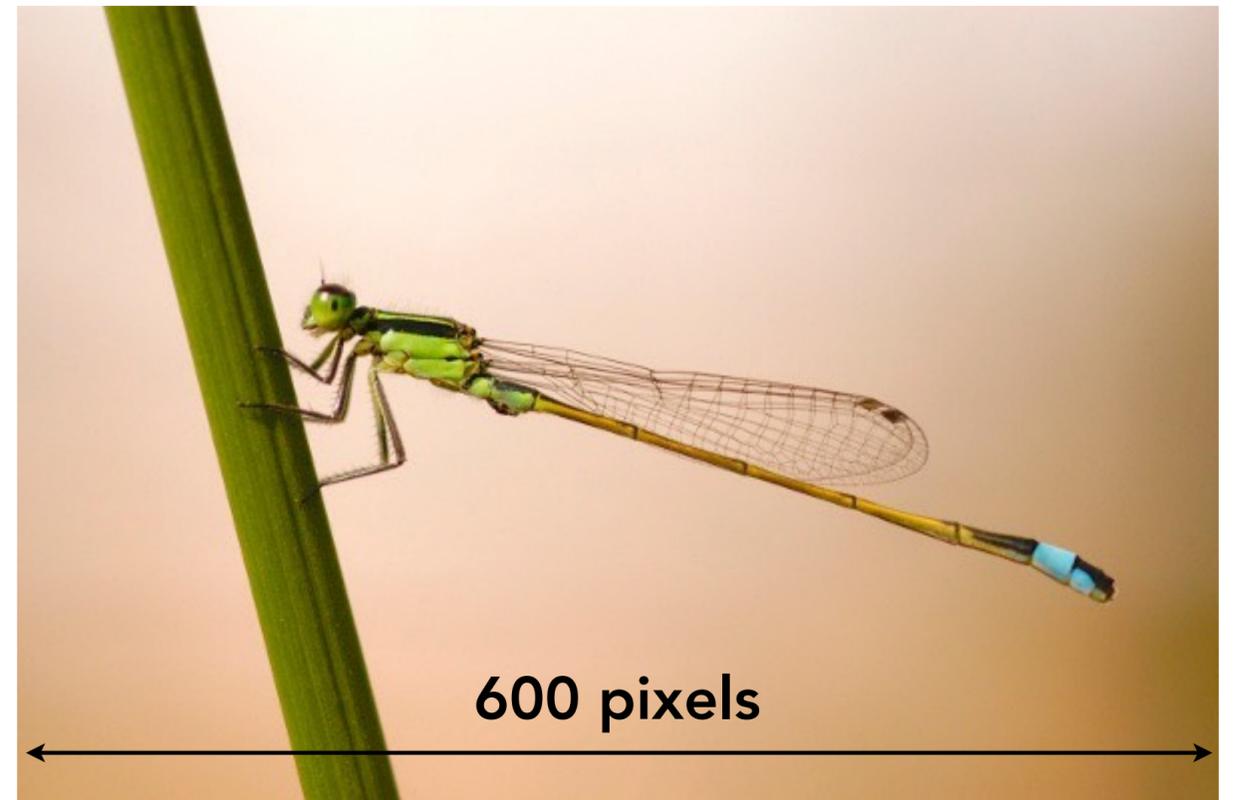
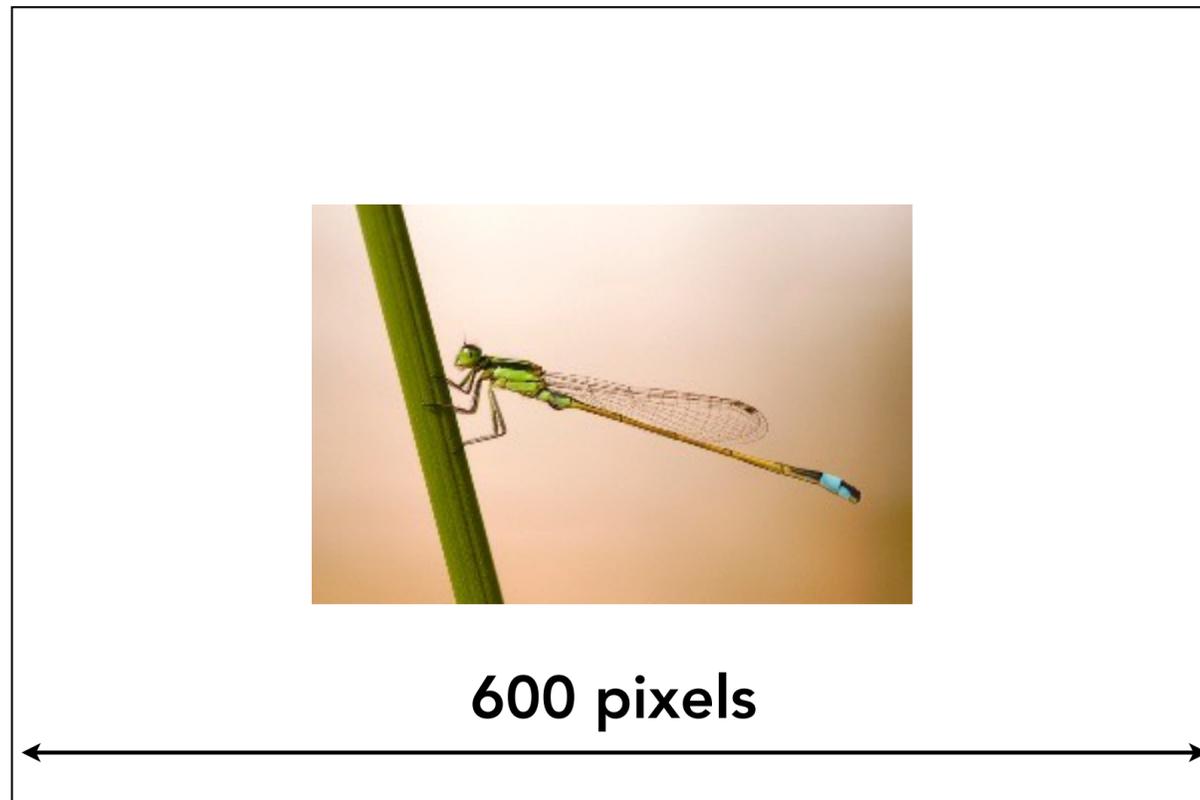


256x256 pixels

Sampling rate on screen vs texture

Rendered image

Texture



Screen space (x,y)

Texture space (u,v)

Red dots = samples needed to render

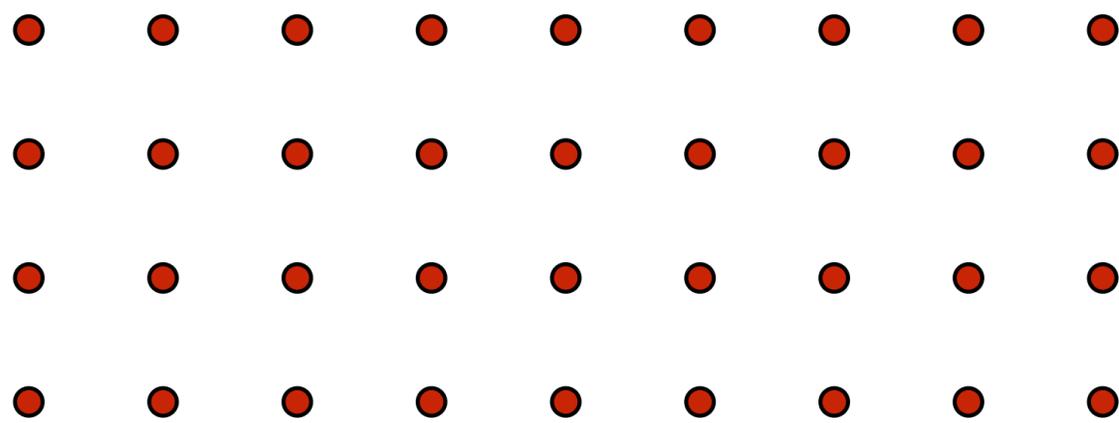
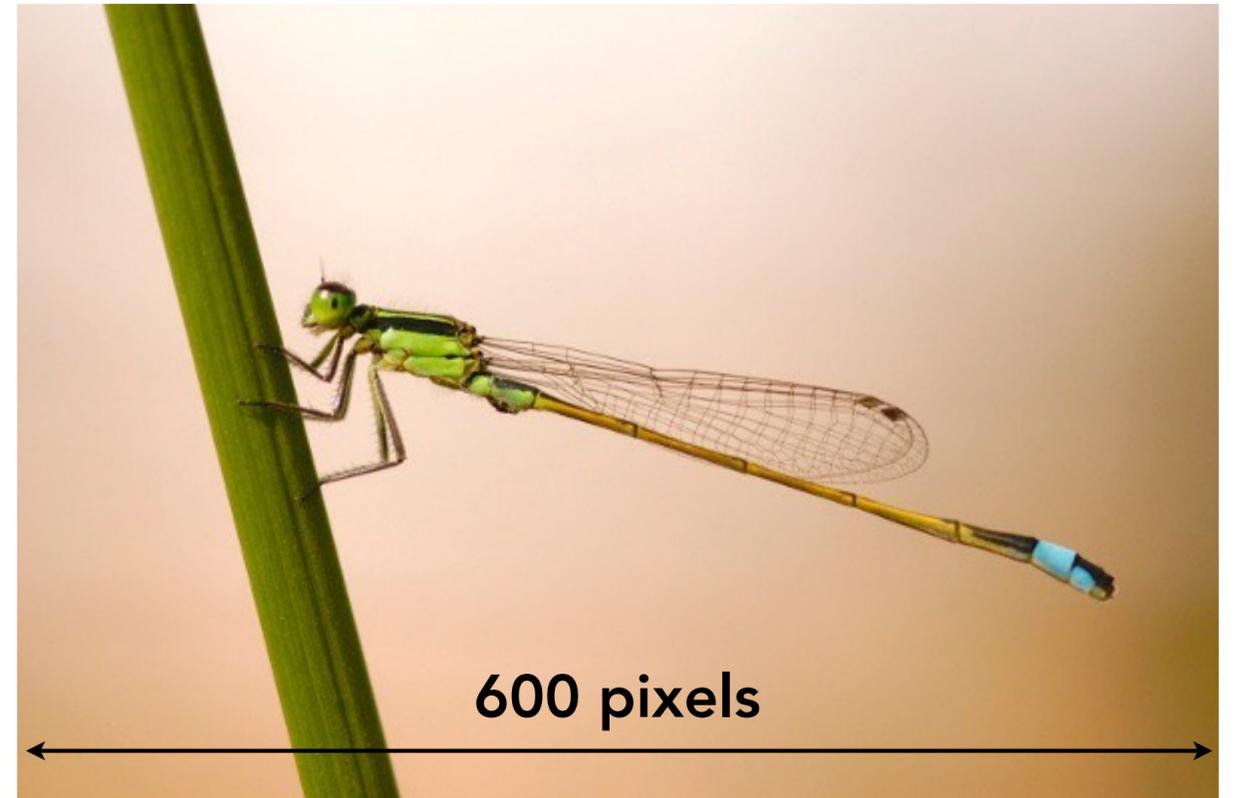
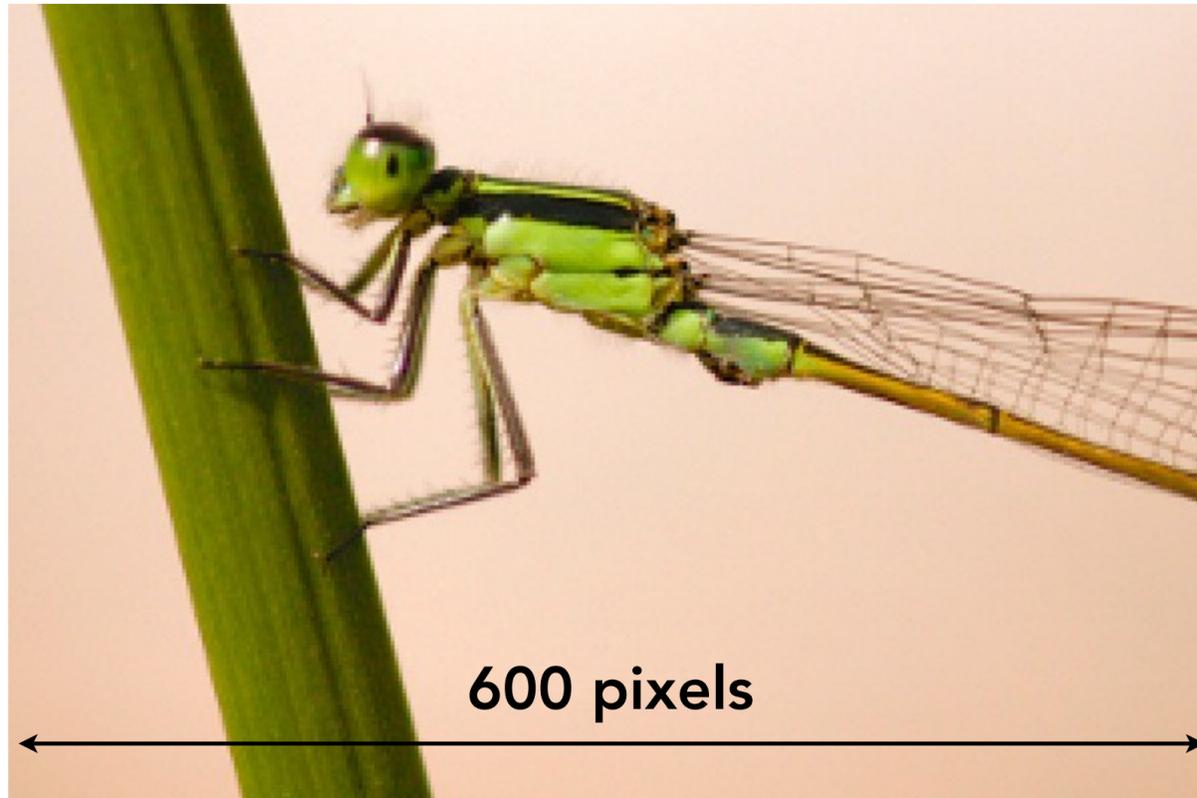
Texture is "minified"

White = samples existing in texture map

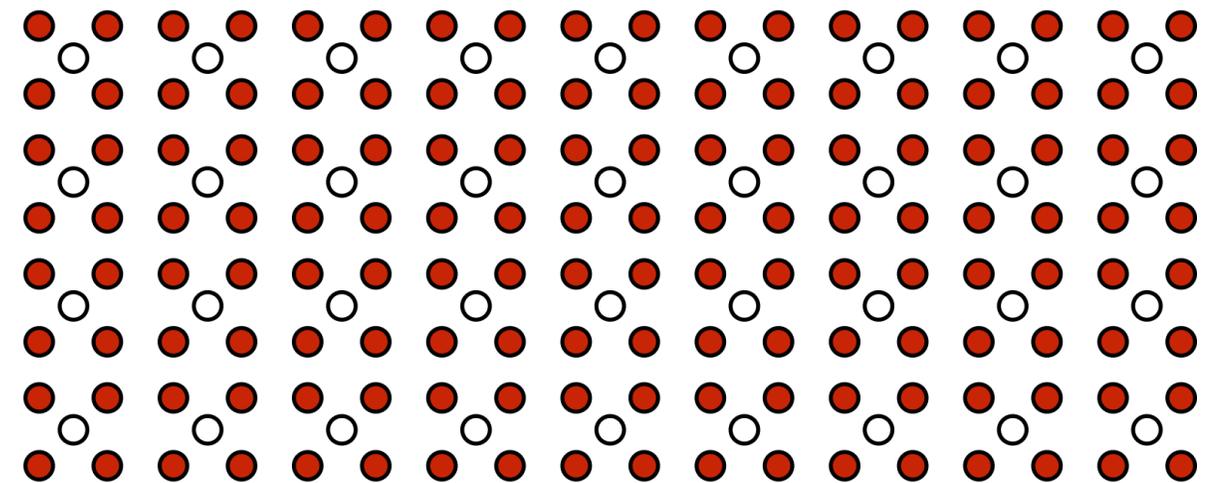
Sampling rate on screen vs texture

Rendered image

Texture



Screen space (x,y)



Texture space (u,v)

Red dots = samples needed to render

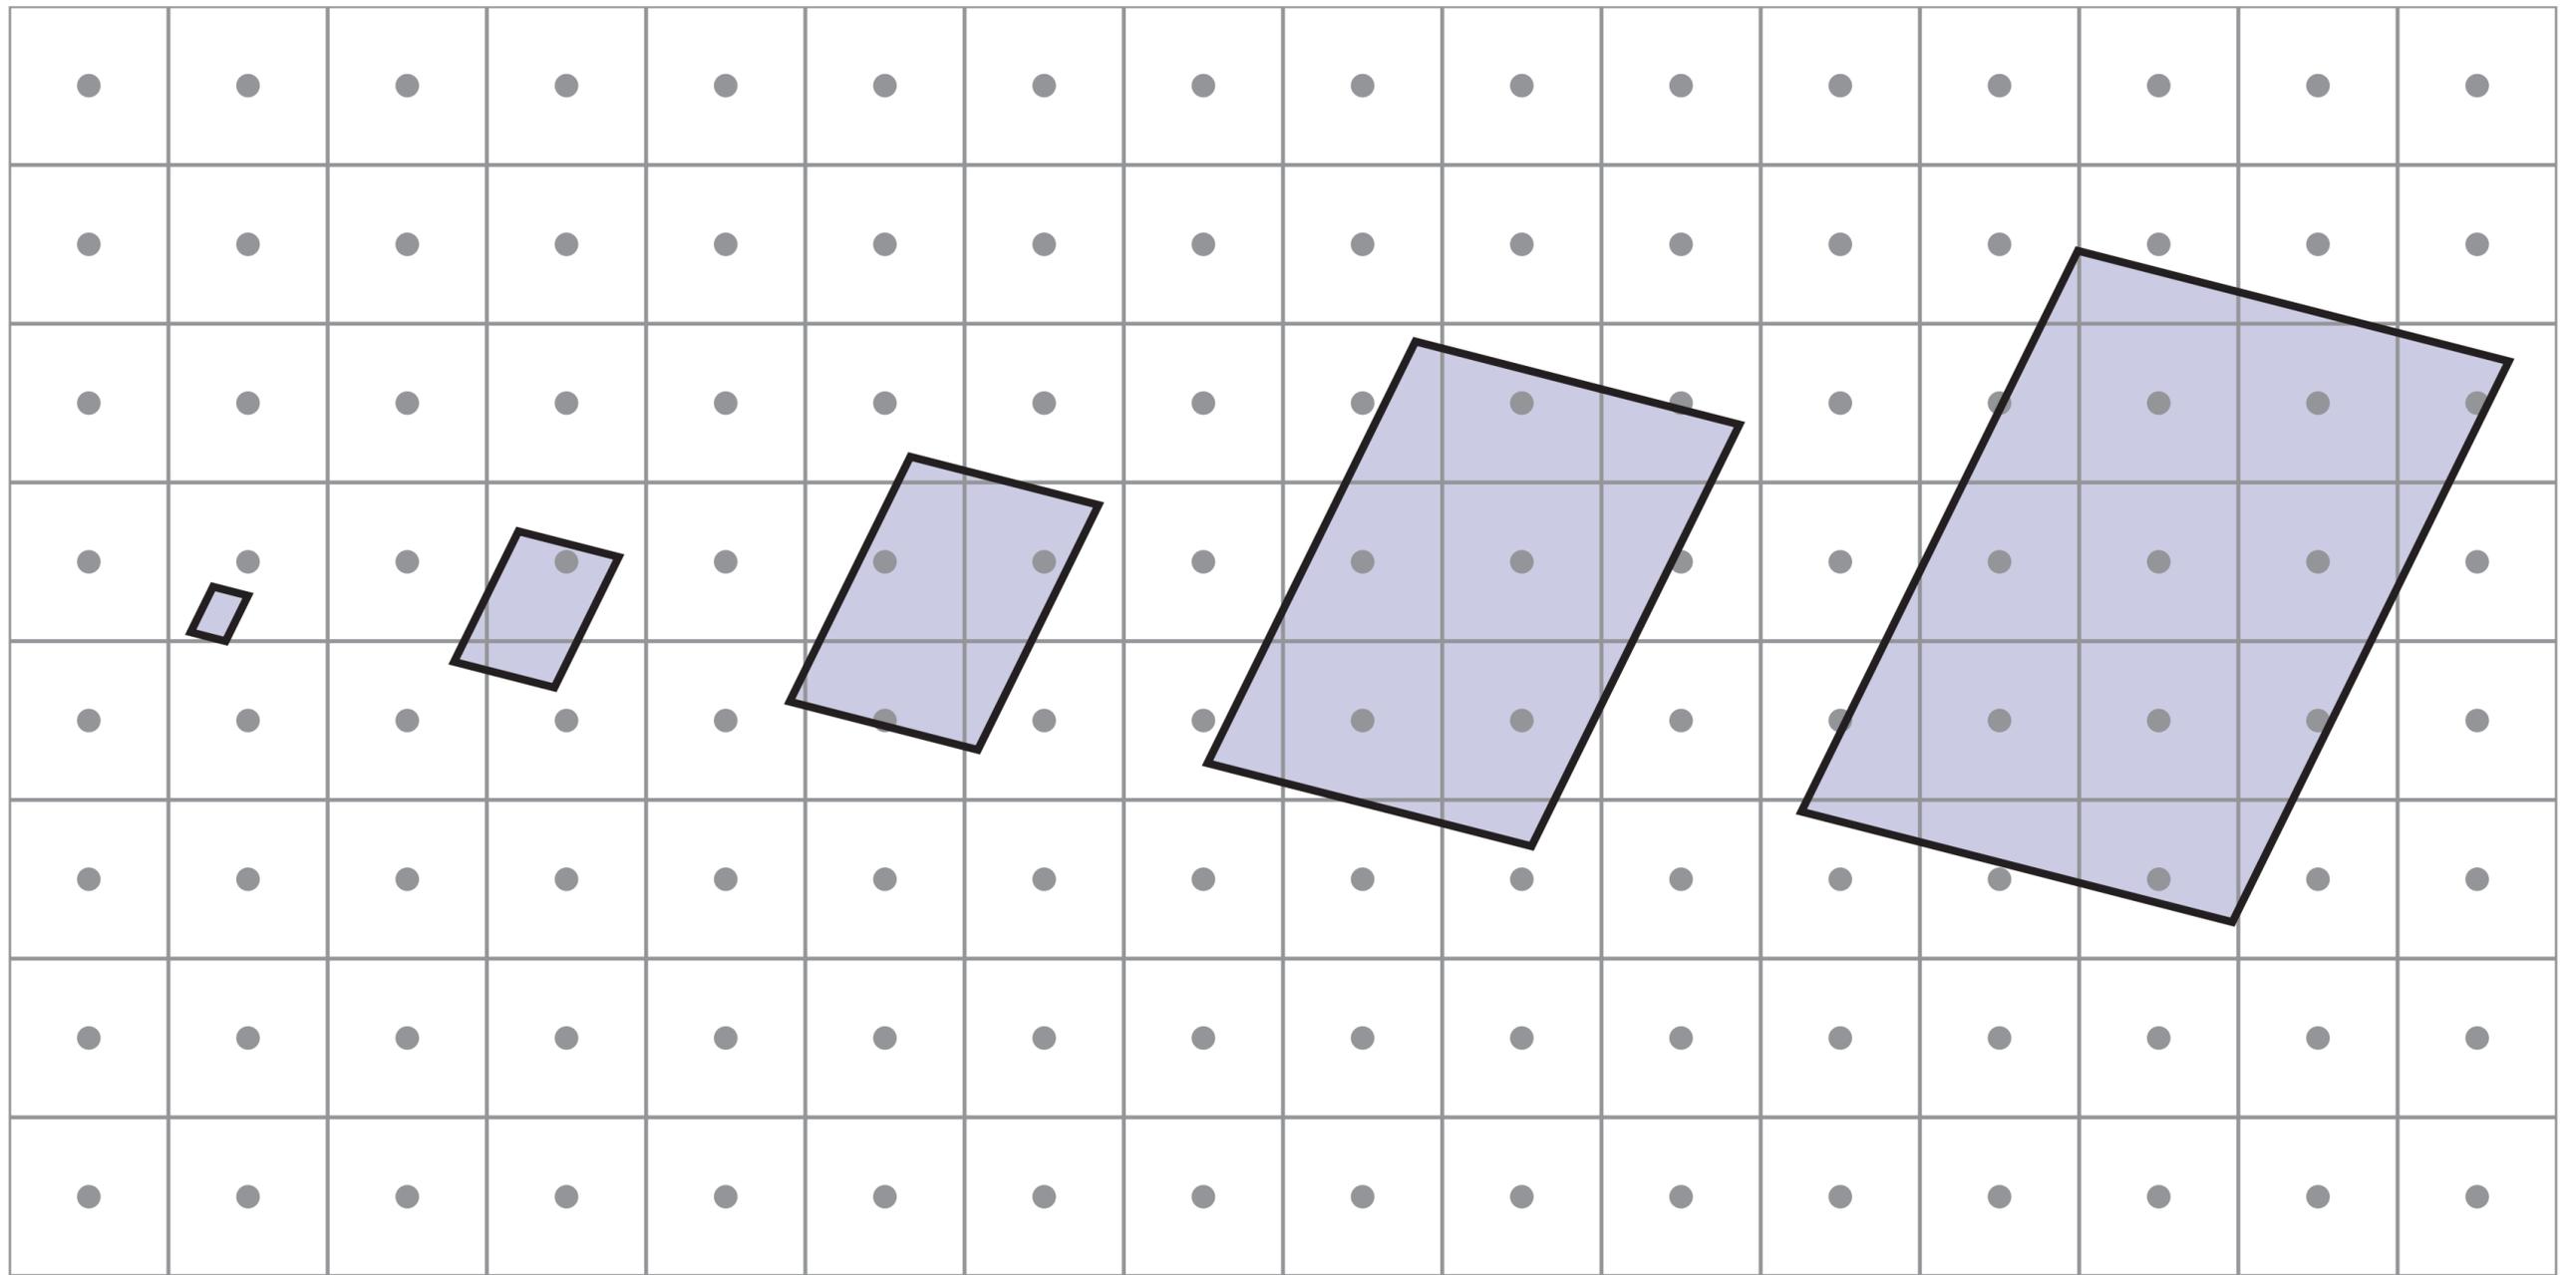
Texture is "magnified"

White = samples existing in texture map

Screen pixel area vs texel area

- **At optimal viewing size:**
 - **1:1 mapping between pixel sampling rate and texel sampling rate**
 - **Dependent on screen and texture resolution! e.g. 512x512**
- **When larger (magnification)**
 - **Multiple pixel samples per texel sample**
- **When smaller (minification)**
 - **One pixel sample per multiple texel samples**

Screen pixel footprint in texture space



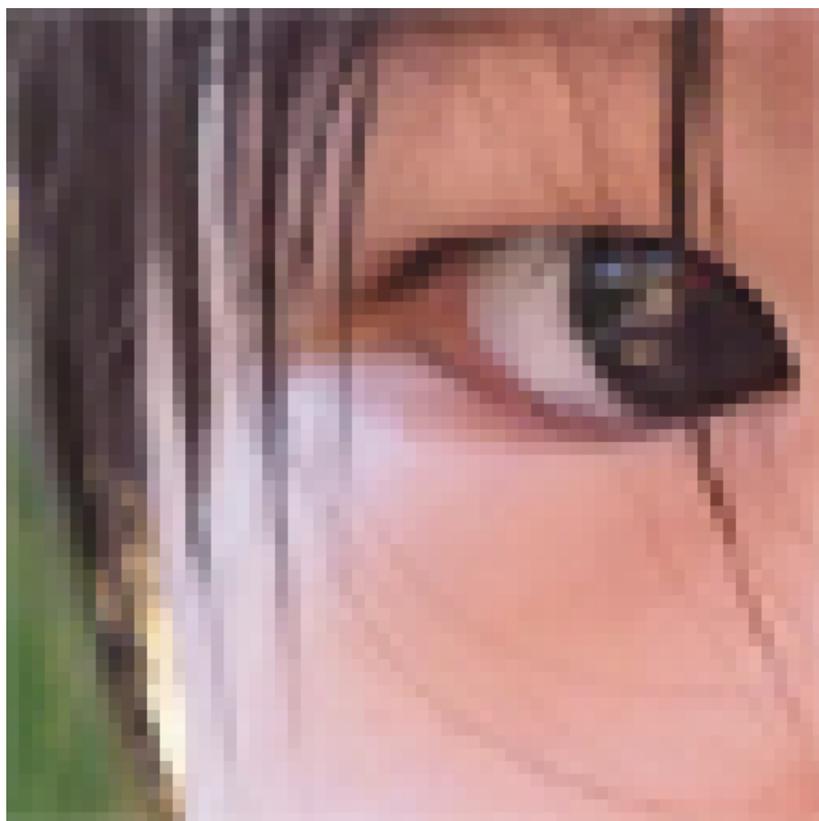
**Upsampling
(Magnification)**



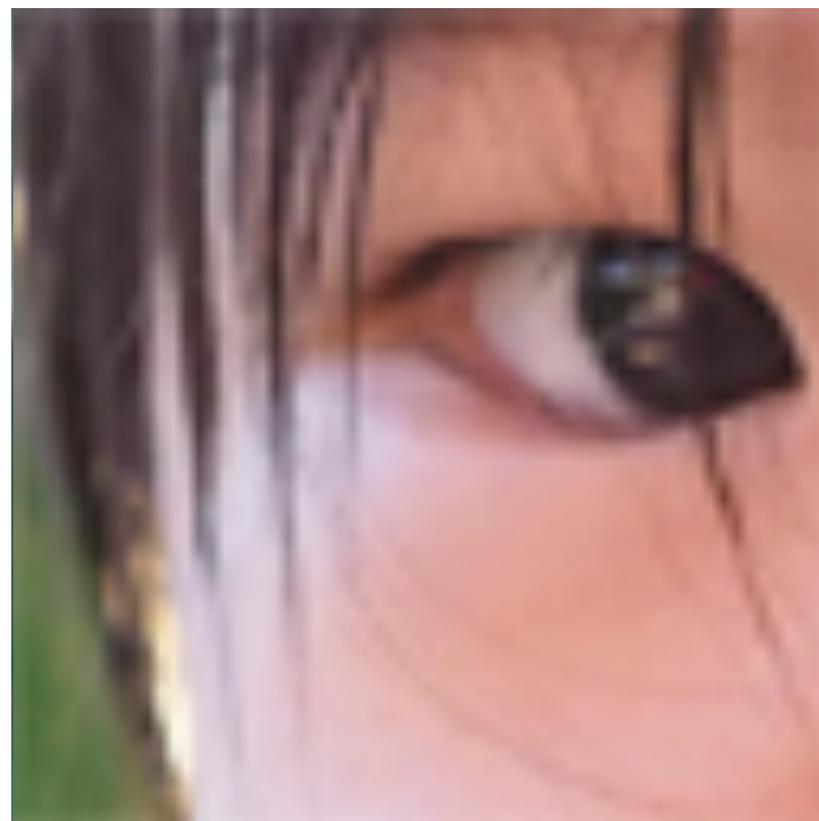
**Downsampling
(Minification)**

Texture magnification - easy case

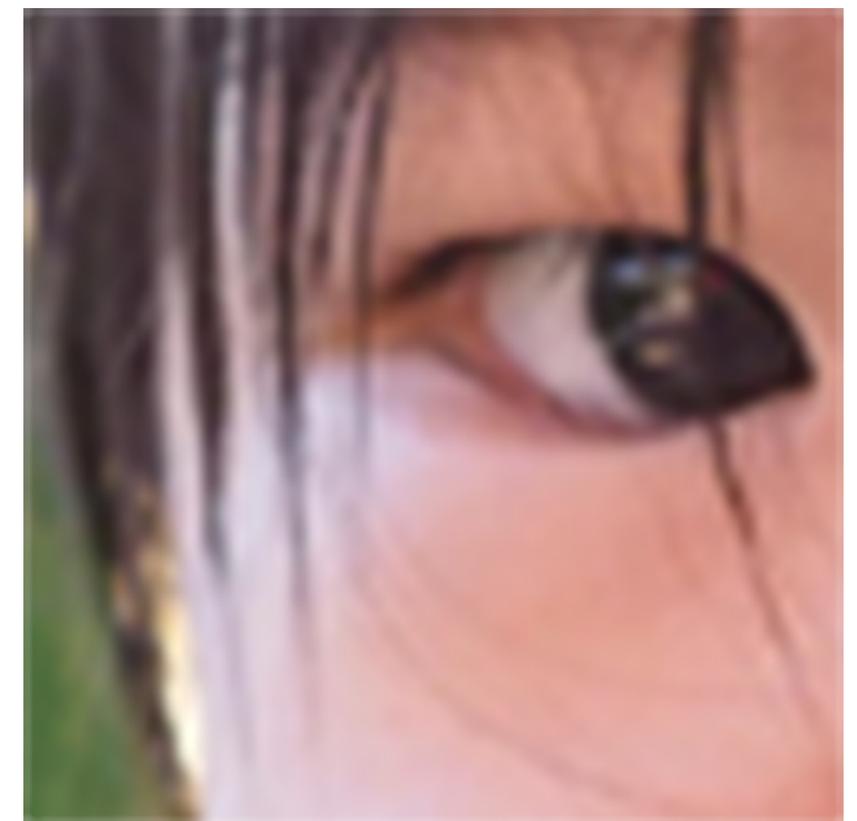
- (Generally don't want this — insufficient texture resolution)
- This is image interpolation (will see kernel function)



Nearest

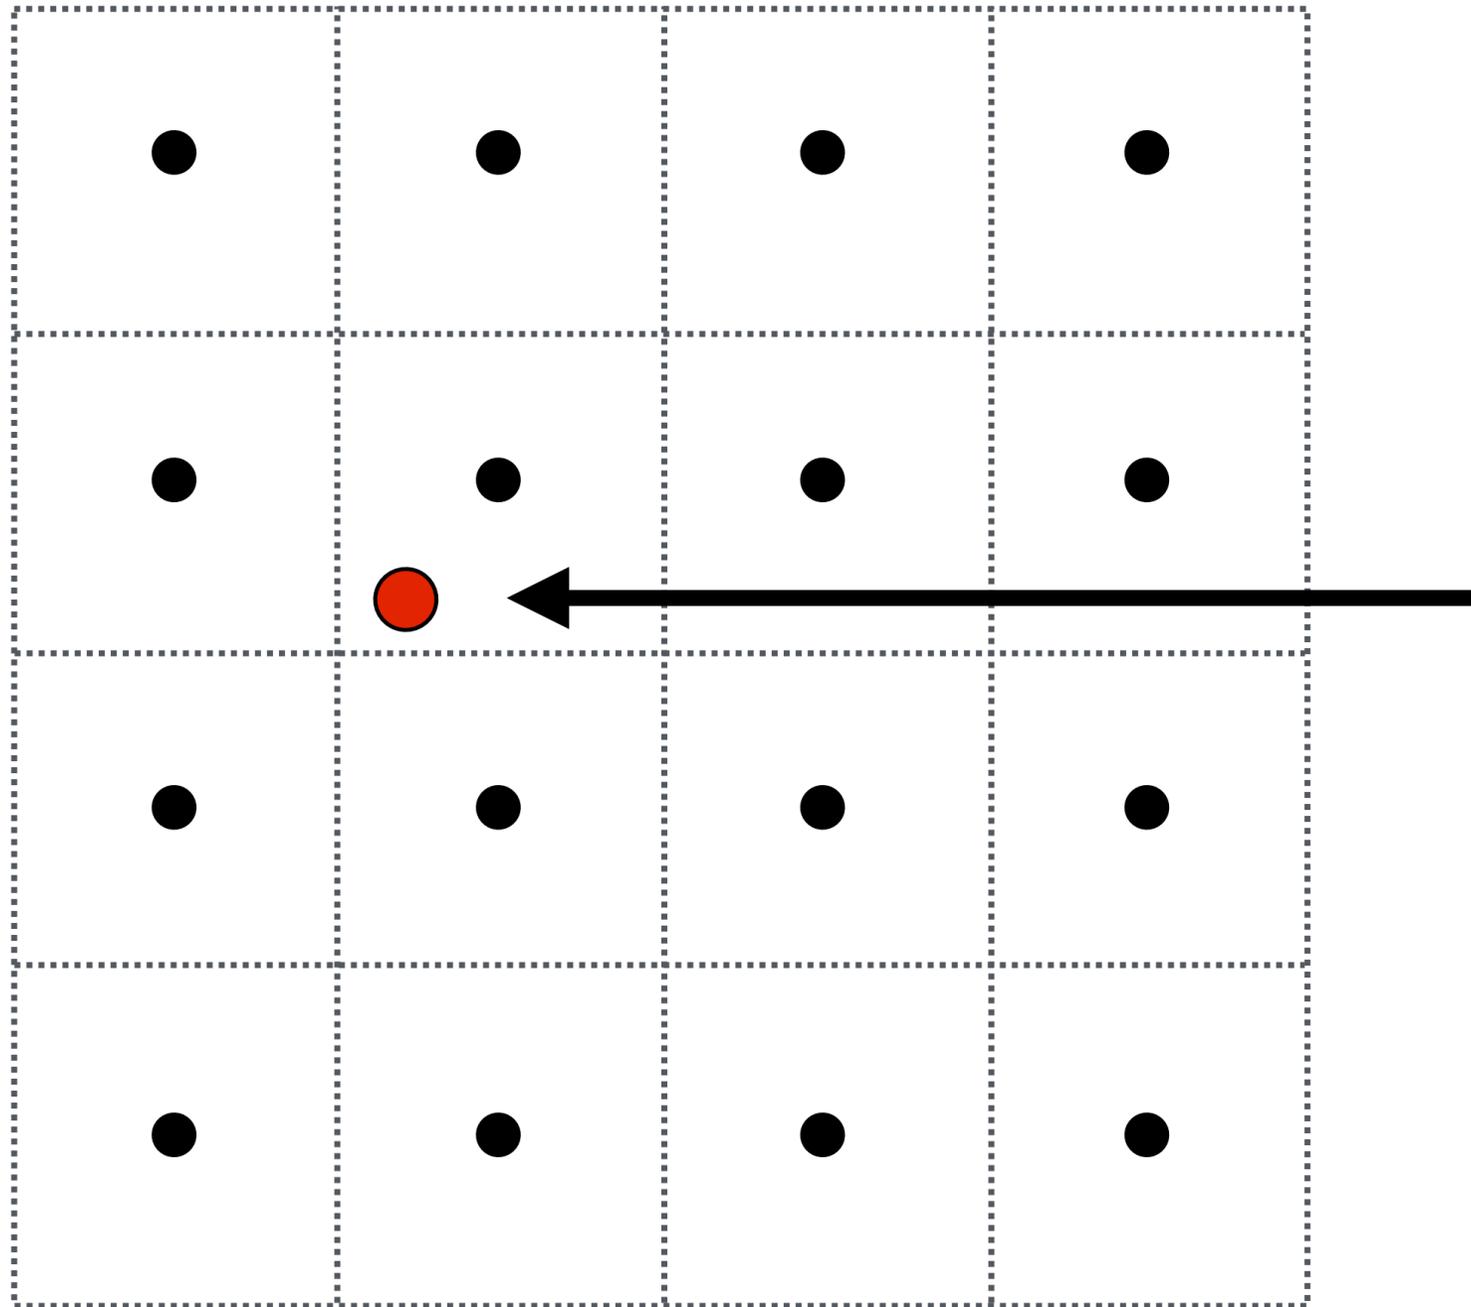


Bilinear



Bicubic

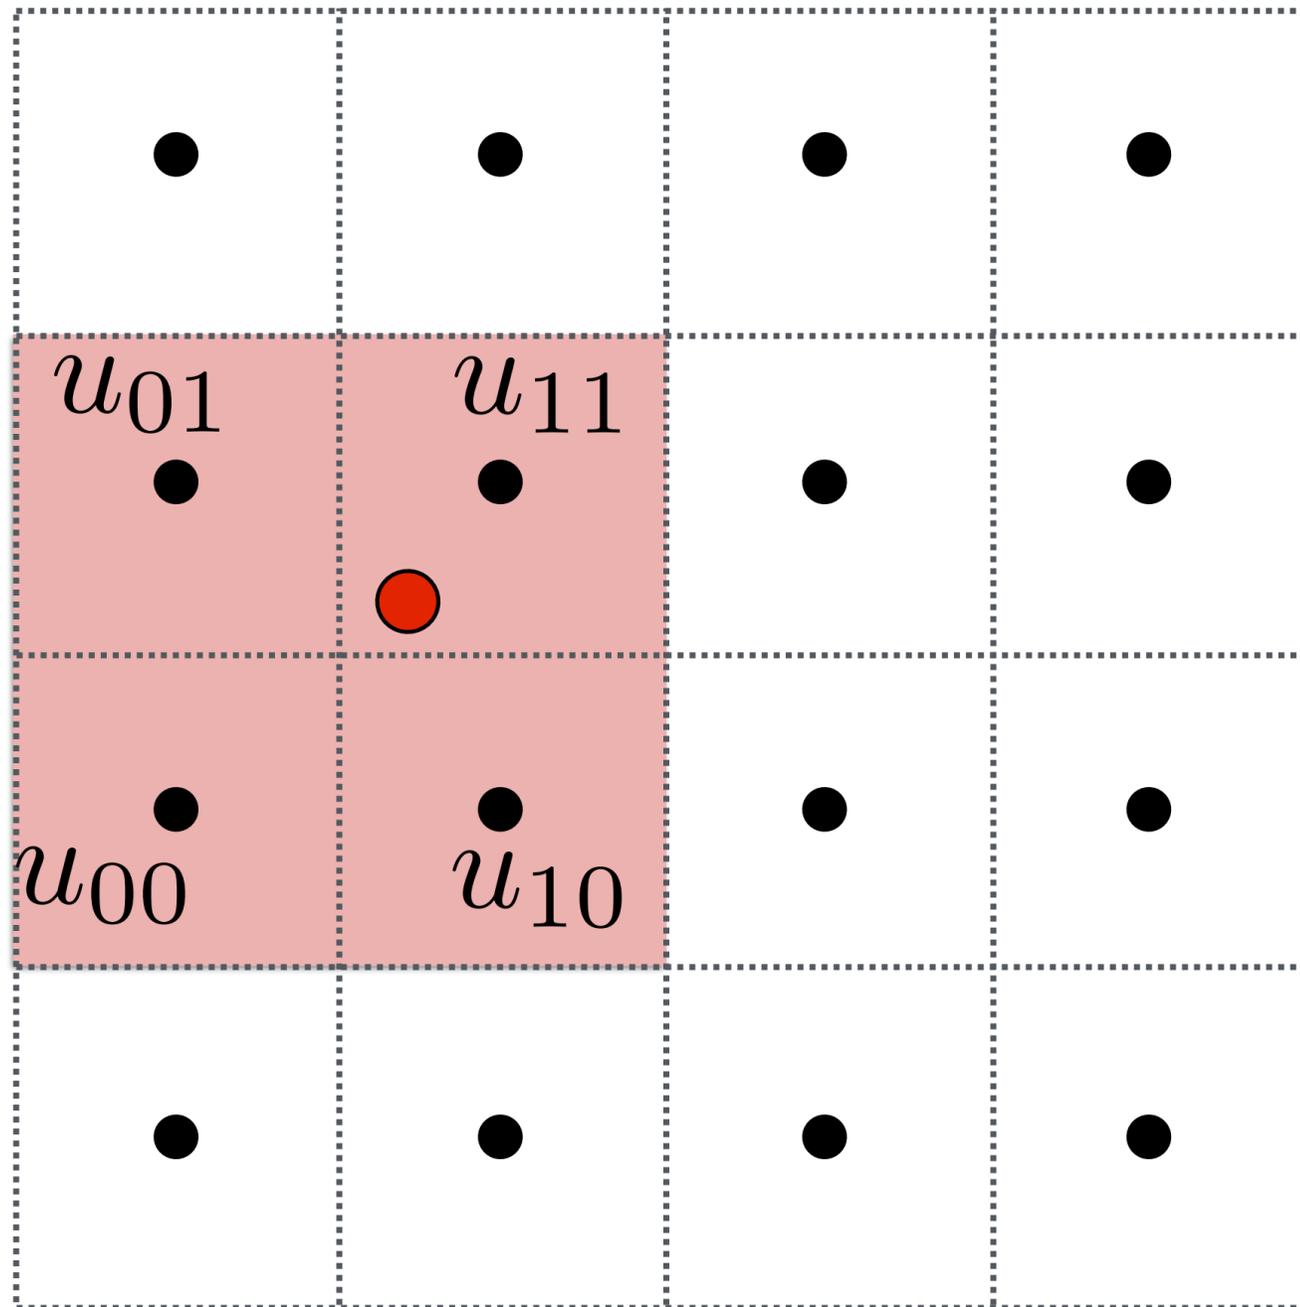
Bilinear Filtering



Want to sample
texture value $f(x,y)$ at
red point

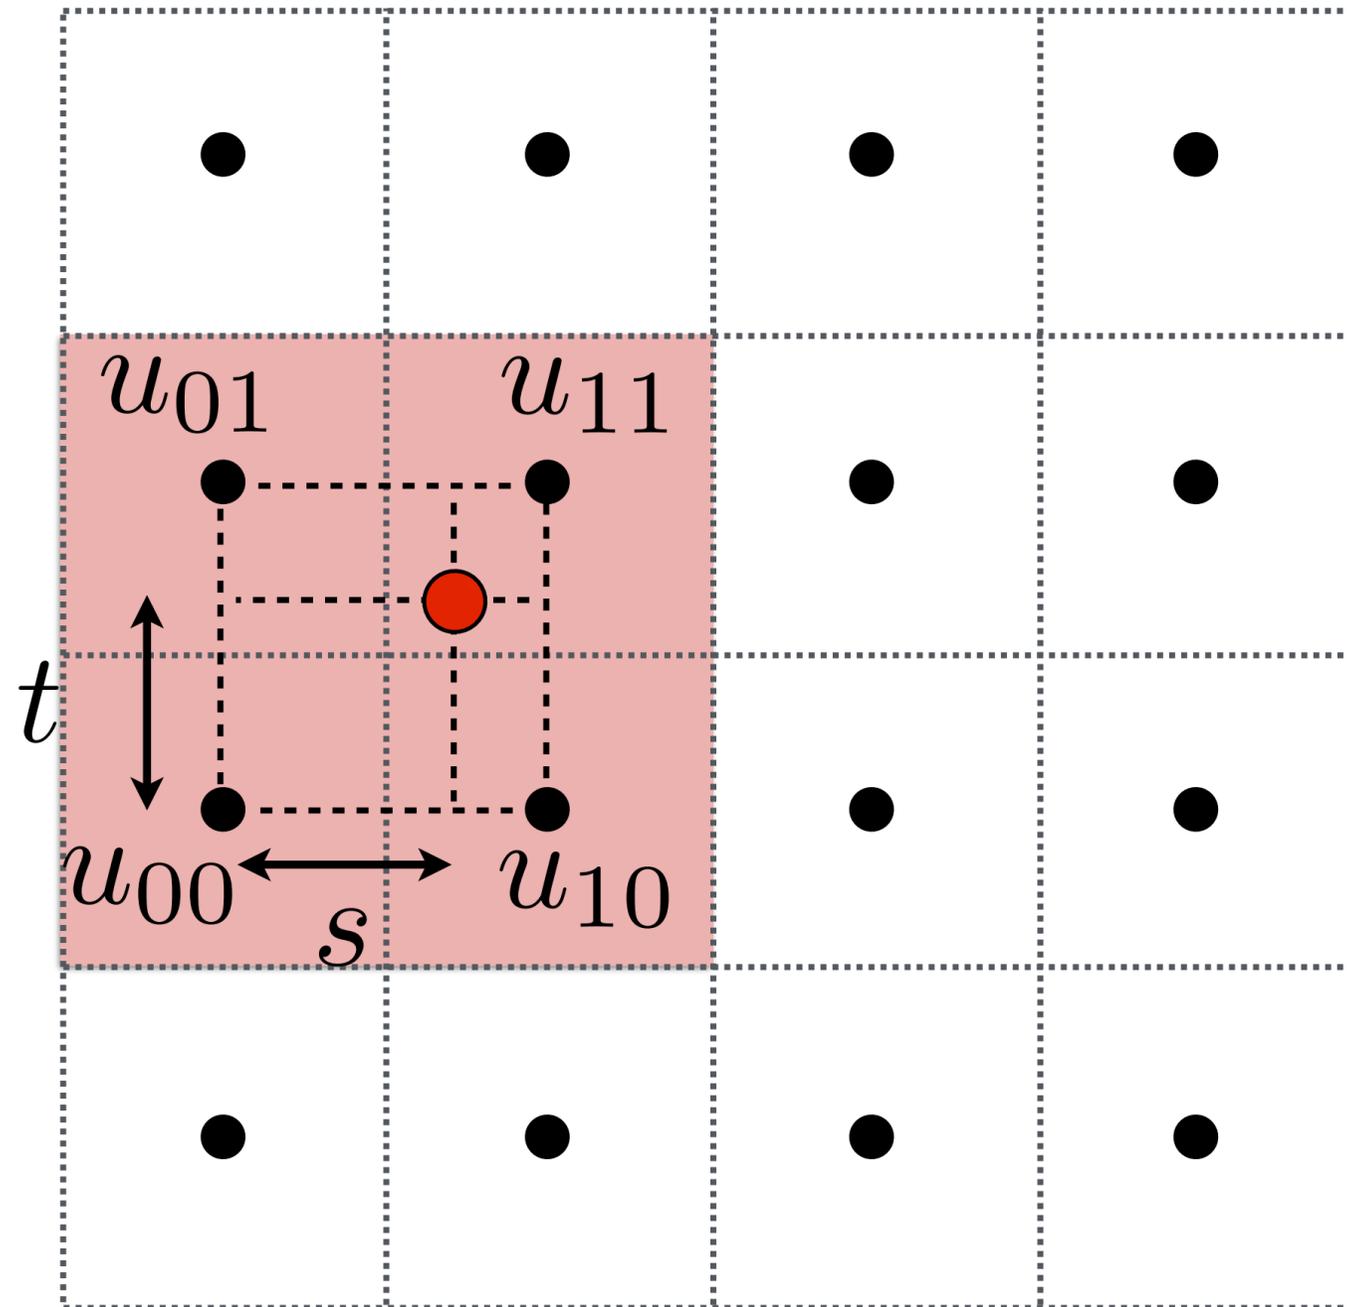
Black points indicate
texture sample
locations

Bilinear filtering



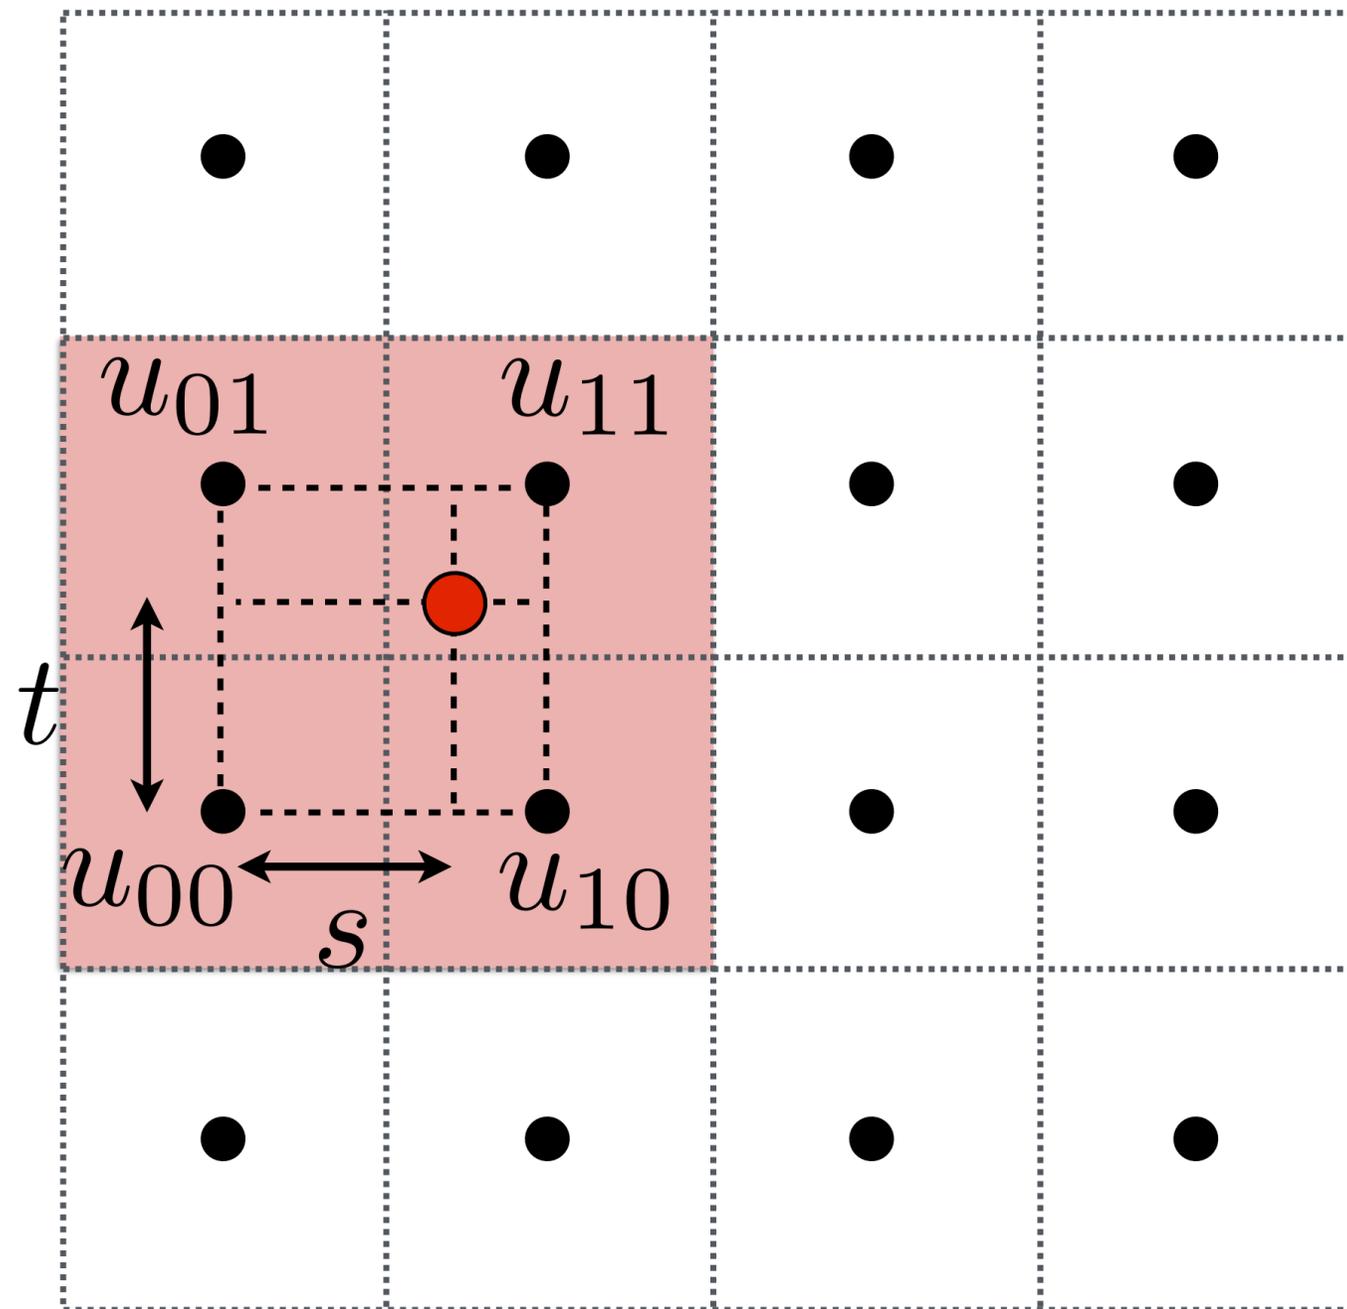
Take 4 nearest sample locations, with texture values as labeled.

Bilinear filtering



And fractional offsets, (s,t) as shown

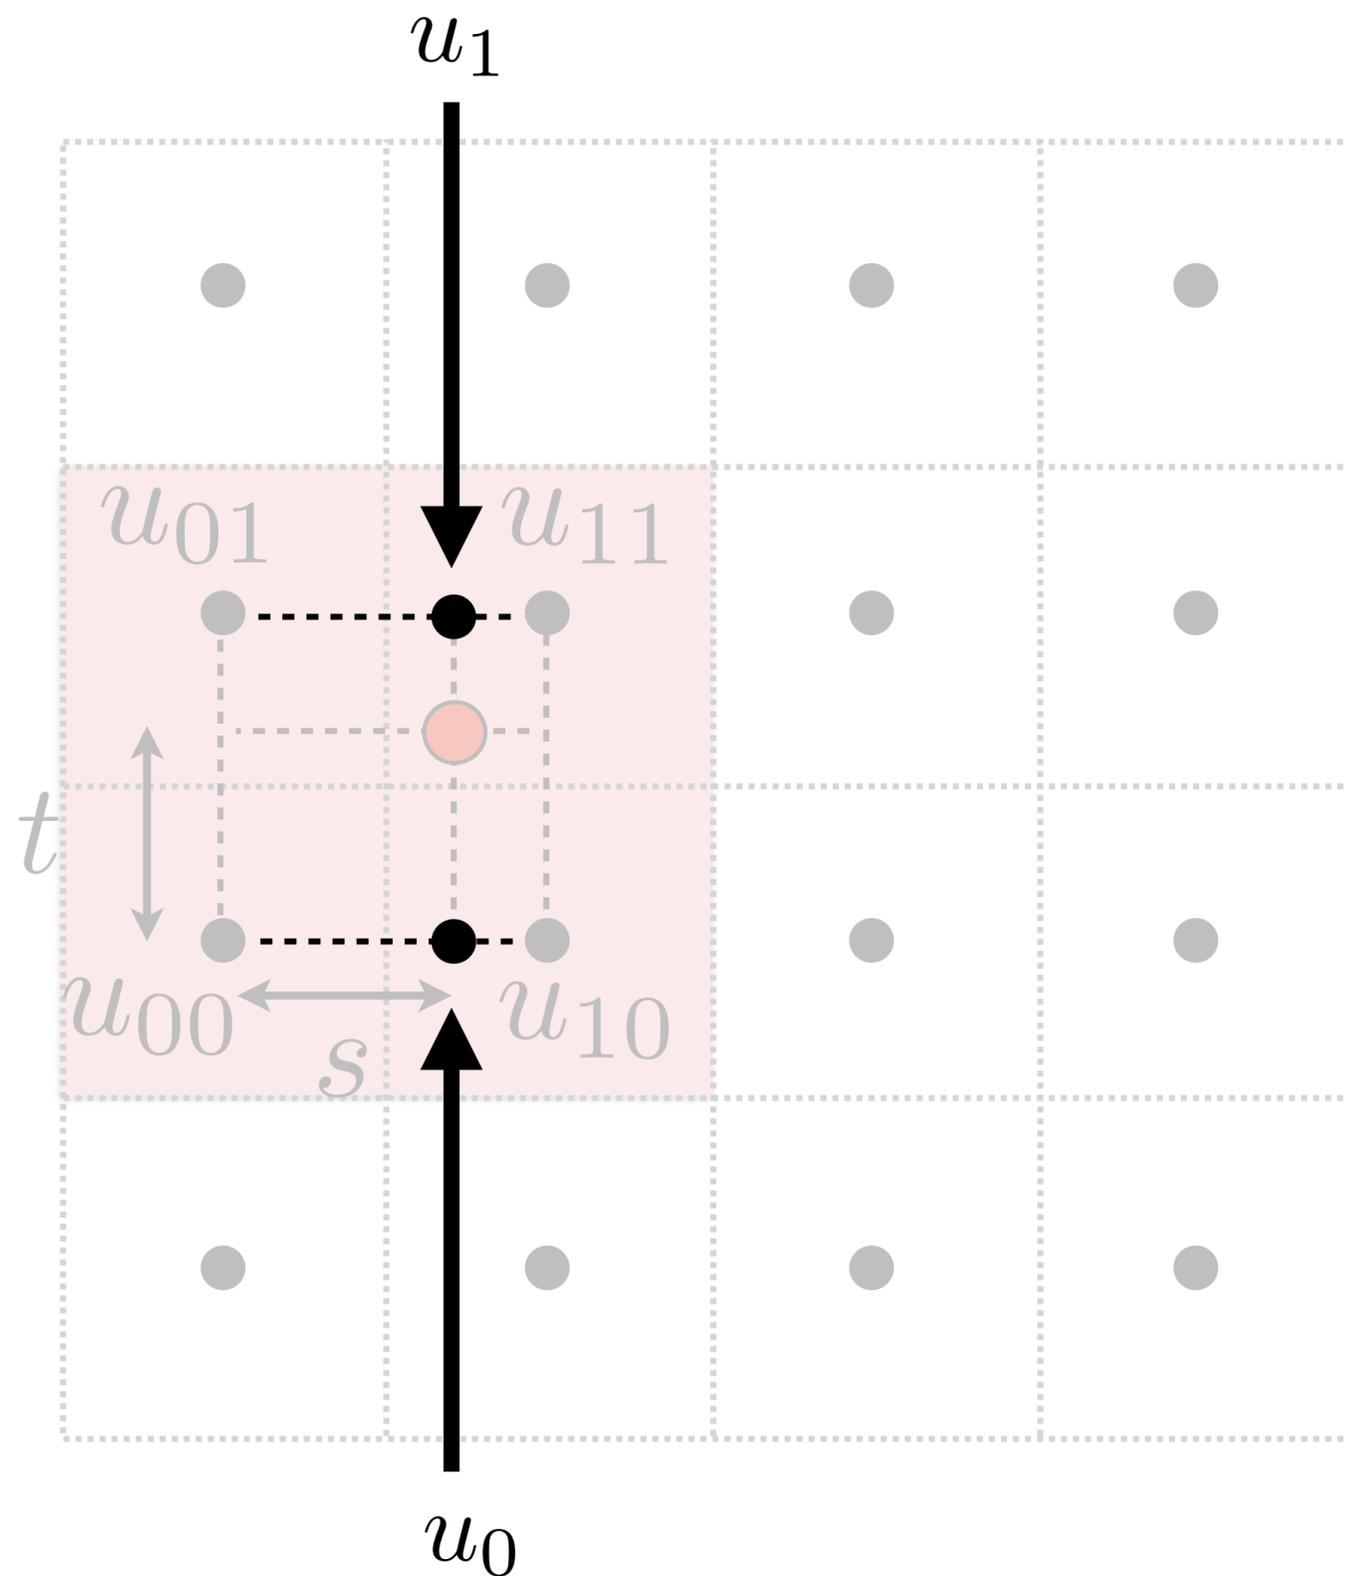
Bilinear filtering



Linear interpolation (1D)

$$\text{lerp}(x, v_0, v_1) = v_0 + x(v_1 - v_0)$$

Bilinear filtering



Linear interpolation (1D)

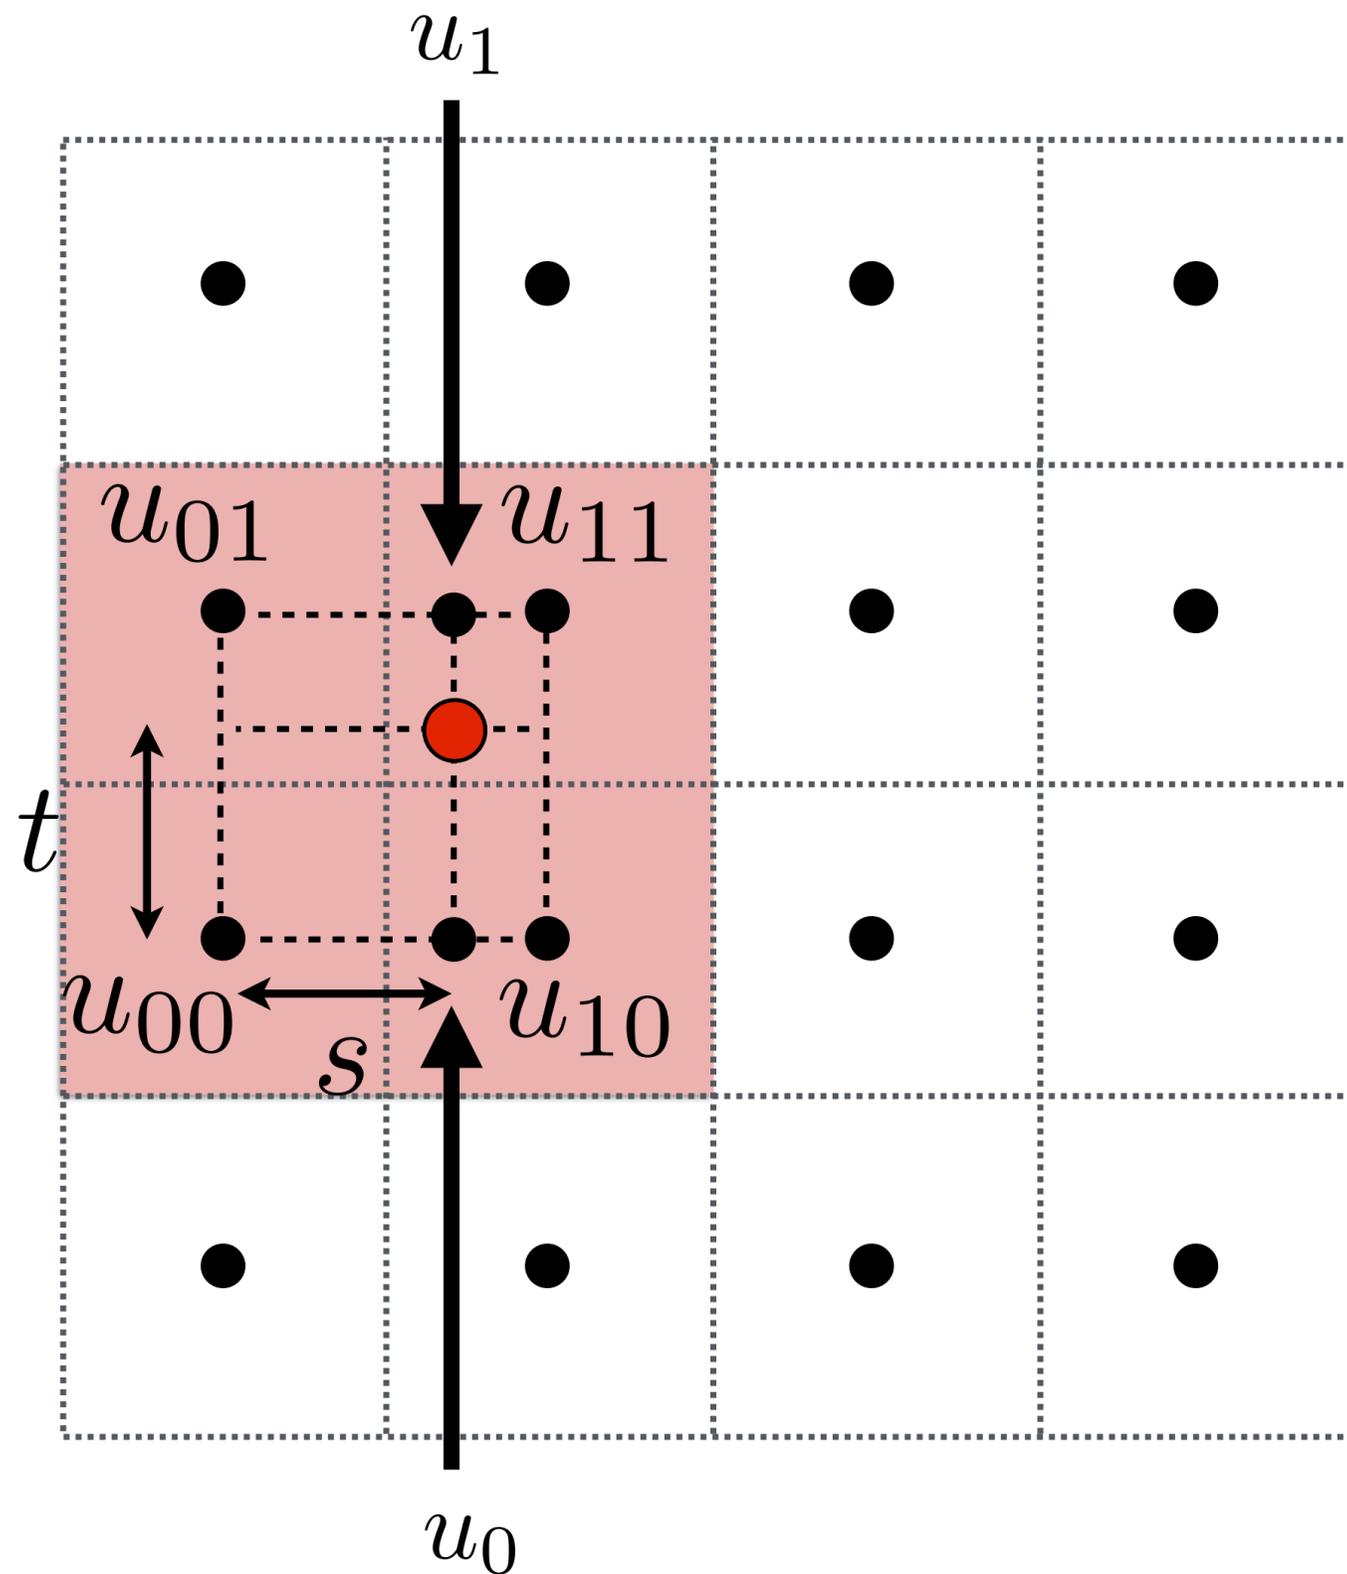
$$\text{lerp}(x, v_0, v_1) = v_0 + x(v_1 - v_0)$$

Two helper lerps (horizontal)

$$u_0 = \text{lerp}(s, u_{00}, u_{10})$$

$$u_1 = \text{lerp}(s, u_{01}, u_{11})$$

Bilinear filtering



Linear interpolation (1D)

$$\text{lerp}(x, v_0, v_1) = v_0 + x(v_1 - v_0)$$

Two helper lerps

$$u_0 = \text{lerp}(s, u_{00}, u_{10})$$

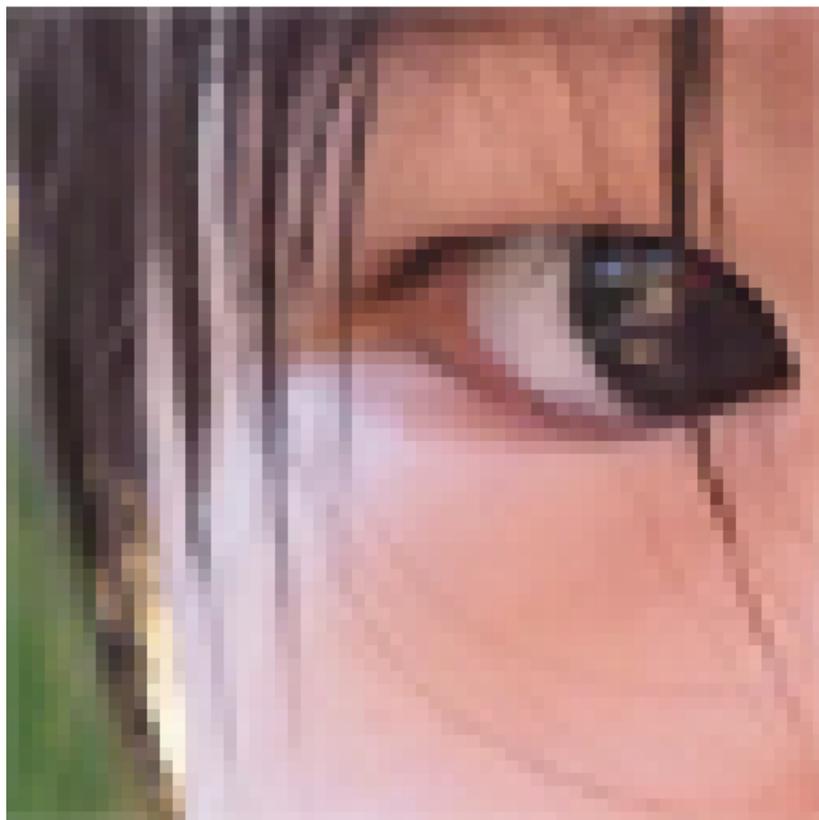
$$u_1 = \text{lerp}(s, u_{01}, u_{11})$$

Final vertical lerp, to get result:

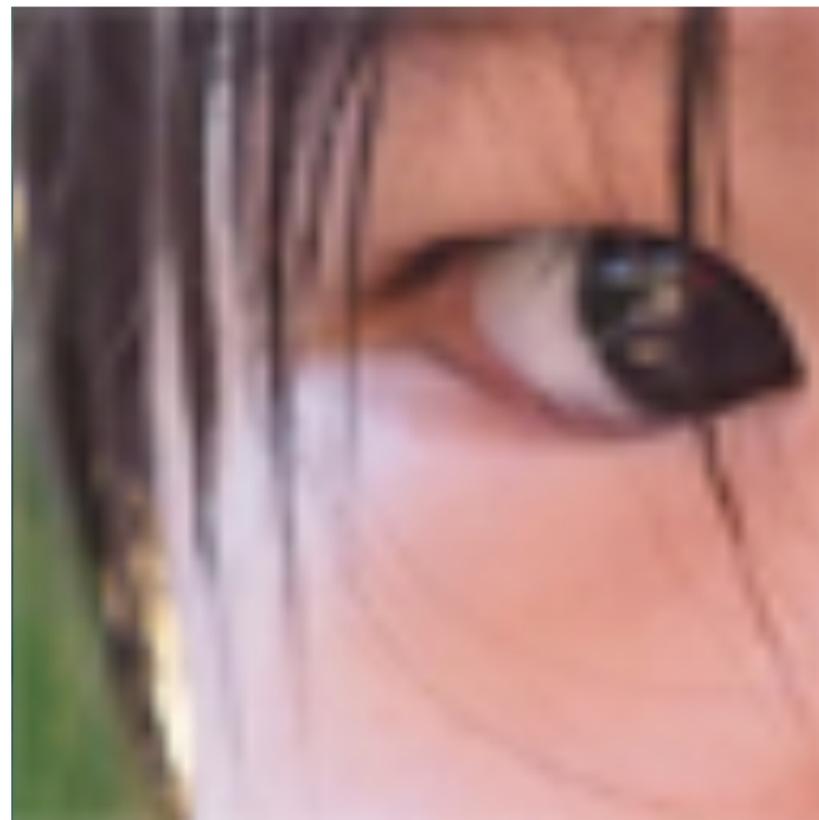
$$f(x, y) = \text{lerp}(t, u_0, u_1)$$

Reconstruction filter function

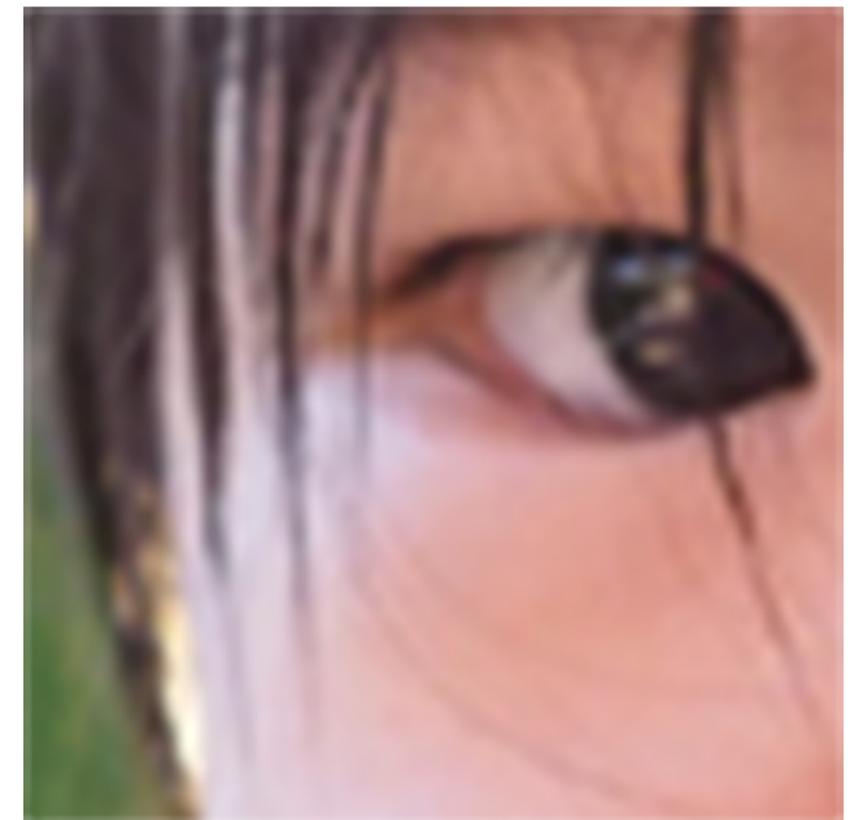
- **Test your understanding:**
 - **What is the reconstruction filter $k(x,y)$ for bilinear interpolation? Nearest? What is a theoretically ideal filter? What are the pros/cons of each?**



Nearest



Bilinear



Bicubic

Texture minification - hard case

■ Challenging

- **Many texels can contribute to pixel footprint**
- **Shape of pixel footprint can be complex**

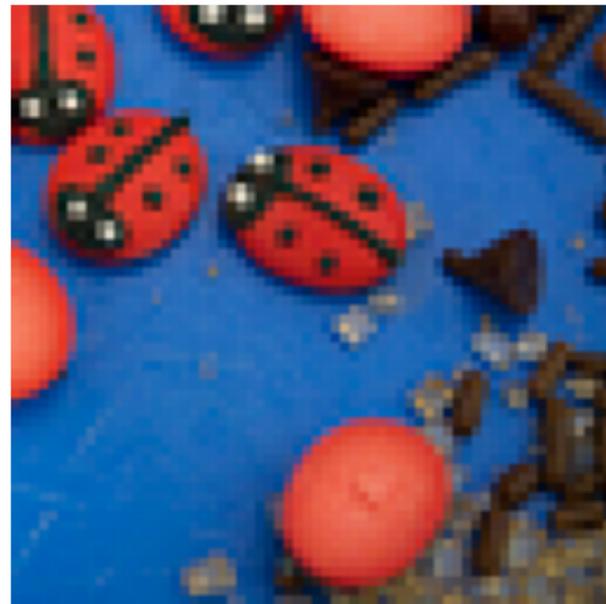
■ Idea:

- **Low-pass filter and downsample texture file, and store successively lower resolutions**
- **For each sample, use the texture file whose resolution approximates the screen sampling rate**

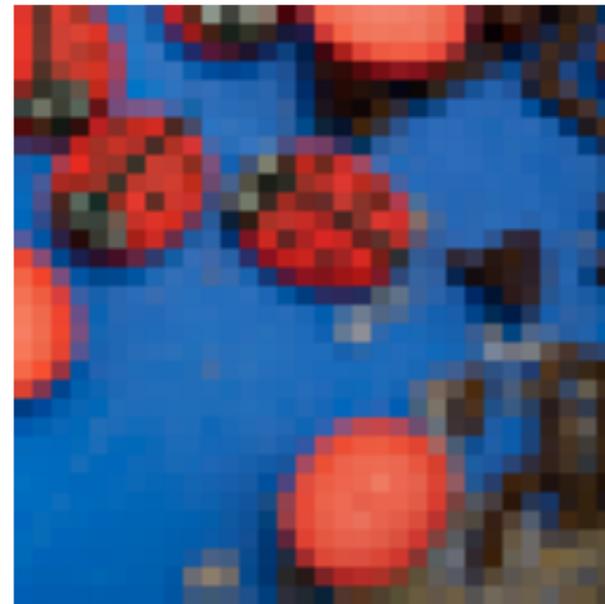
Mipmap (L. Williams 83)



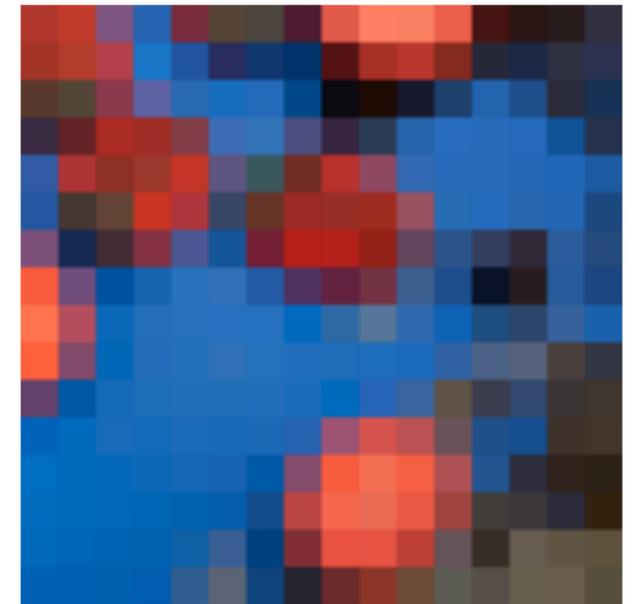
Level 0 = 128x128



Level 1 = 64x64



Level 2 = 32x32



Level 3 = 16x16



Level 4 = 8x8



Level 5 = 4x4



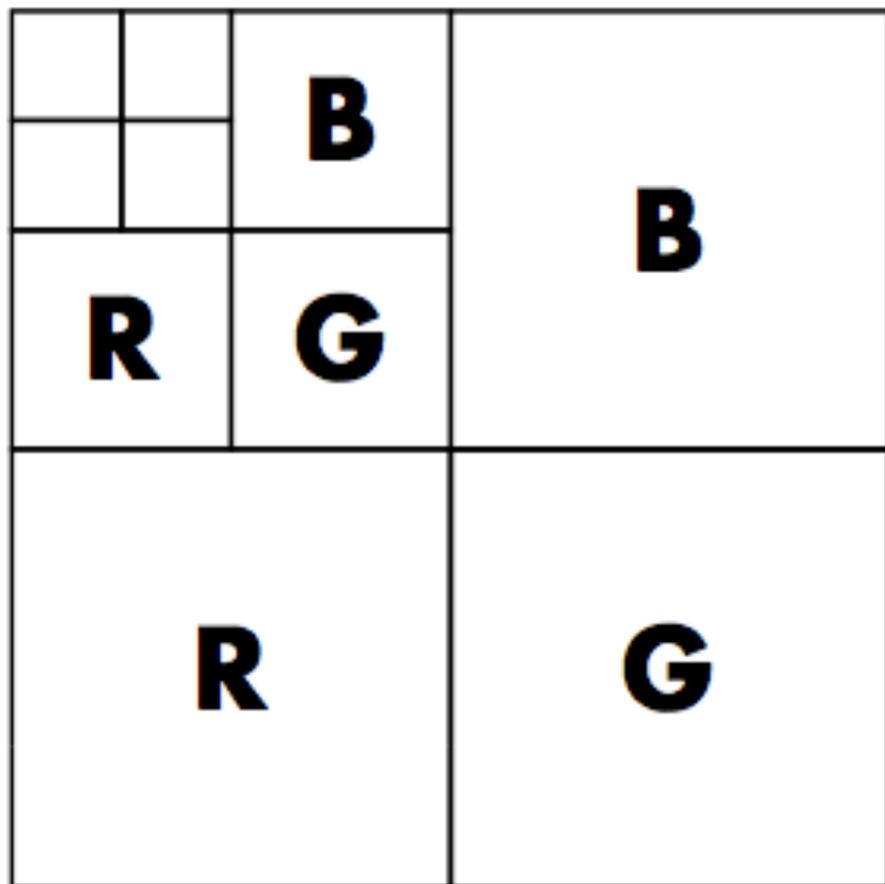
Level 6 = 2x2



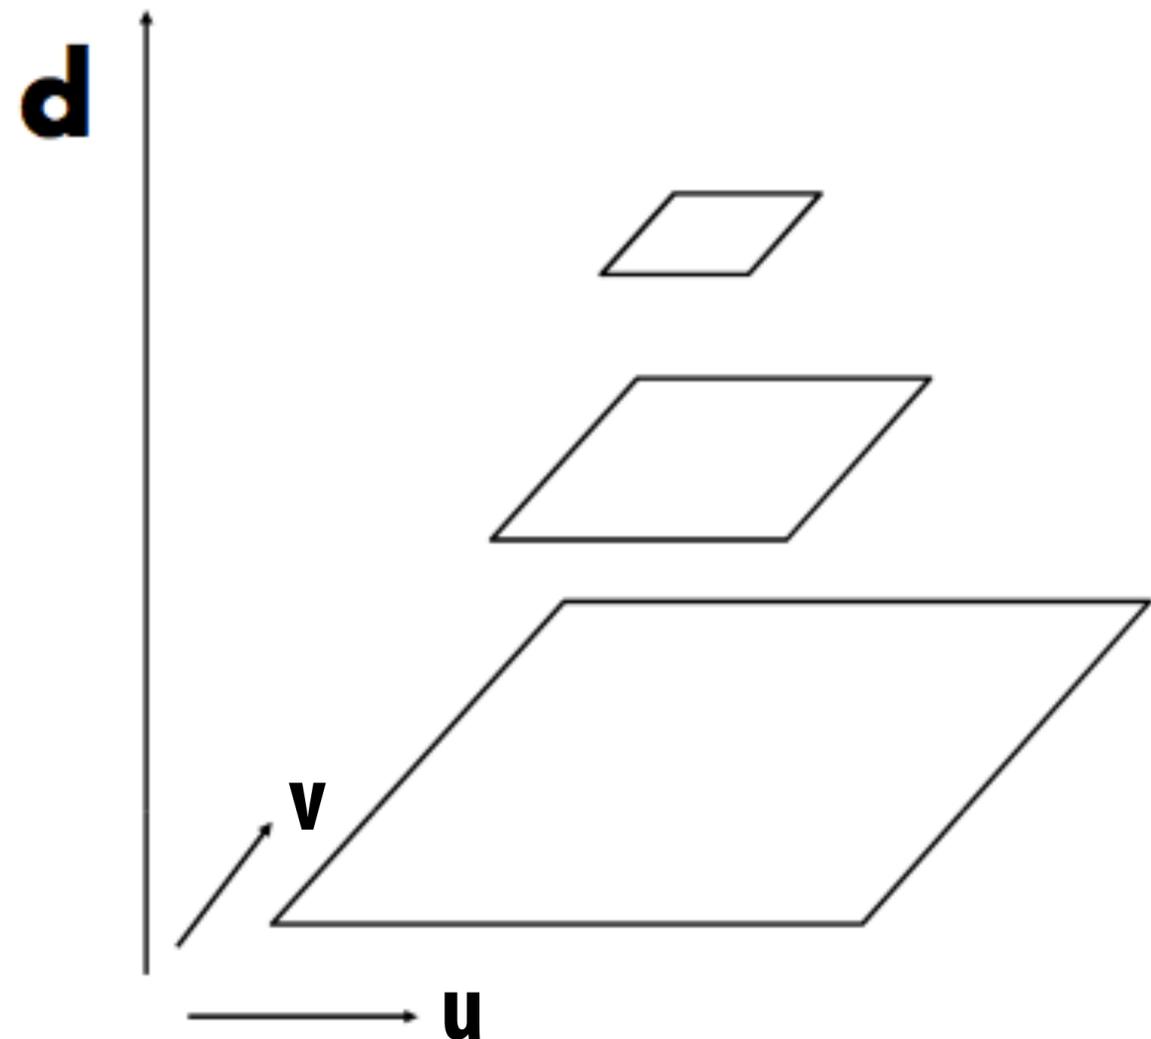
Level 7 = 1x1

"Mip" comes from the Latin "multum in parvo", meaning a multitude in a small space

Mipmap (L. Williams 83)



Williams' original proposed
mip-map layout

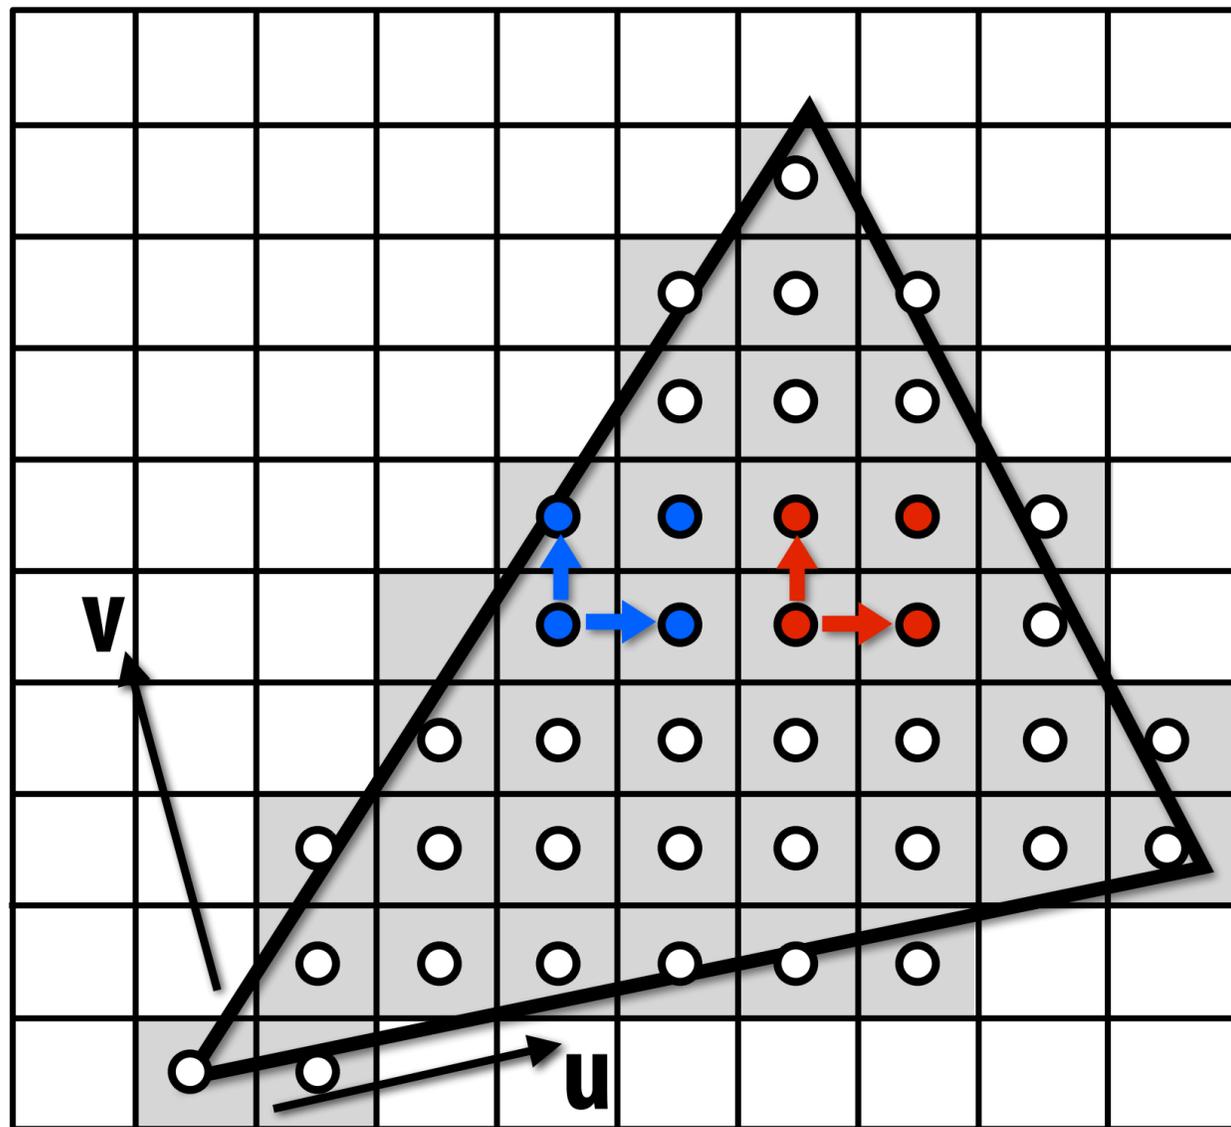


"Mip hierarchy"
level = d

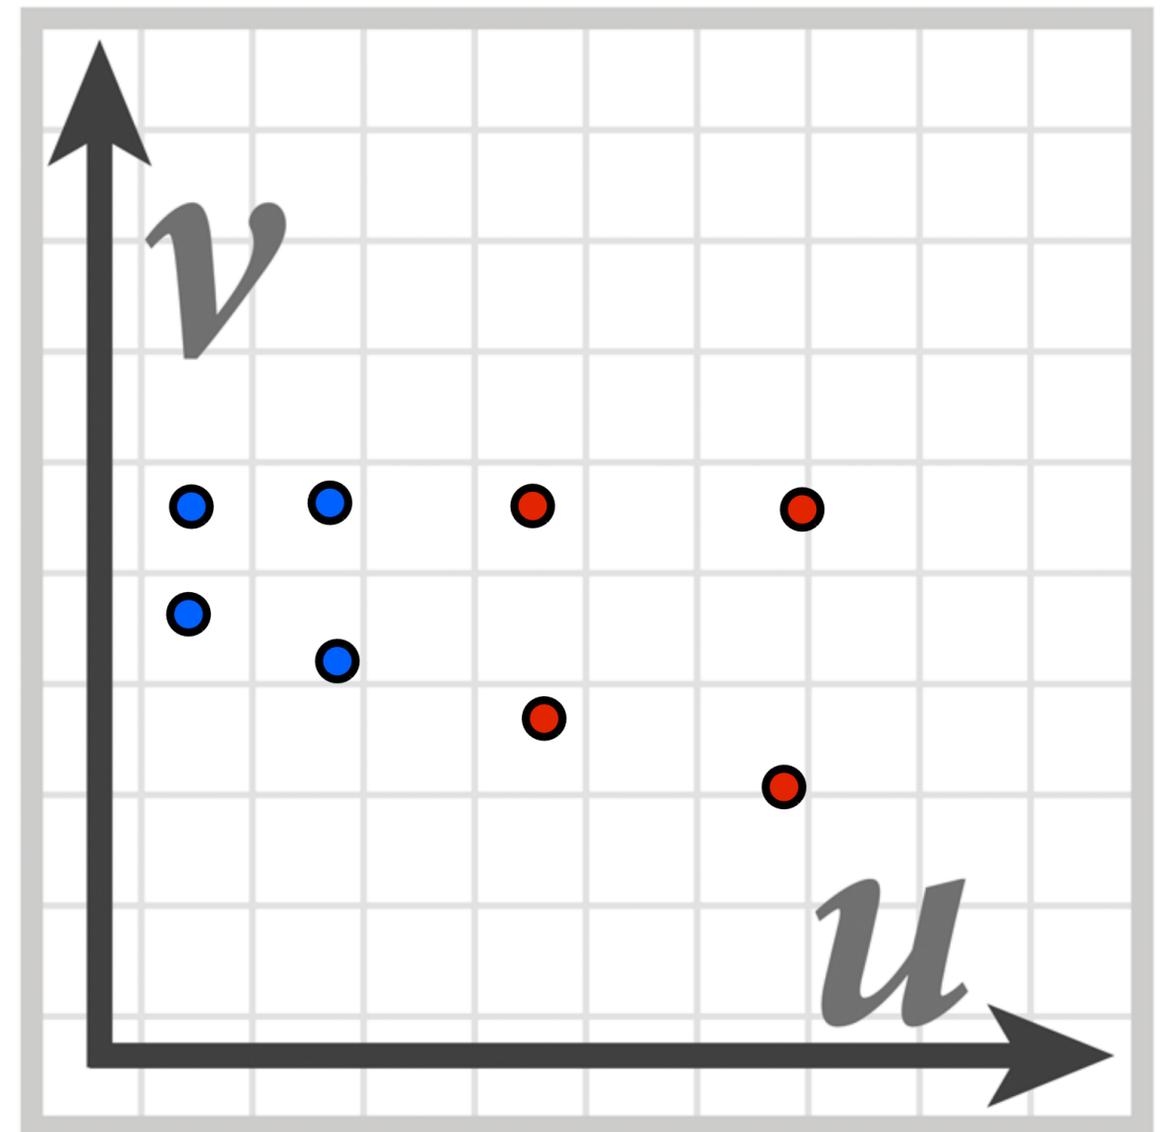
What is the storage overhead of a mipmap?

Computing mipmap level

Compute differences between texture coordinate values of neighboring screen samples



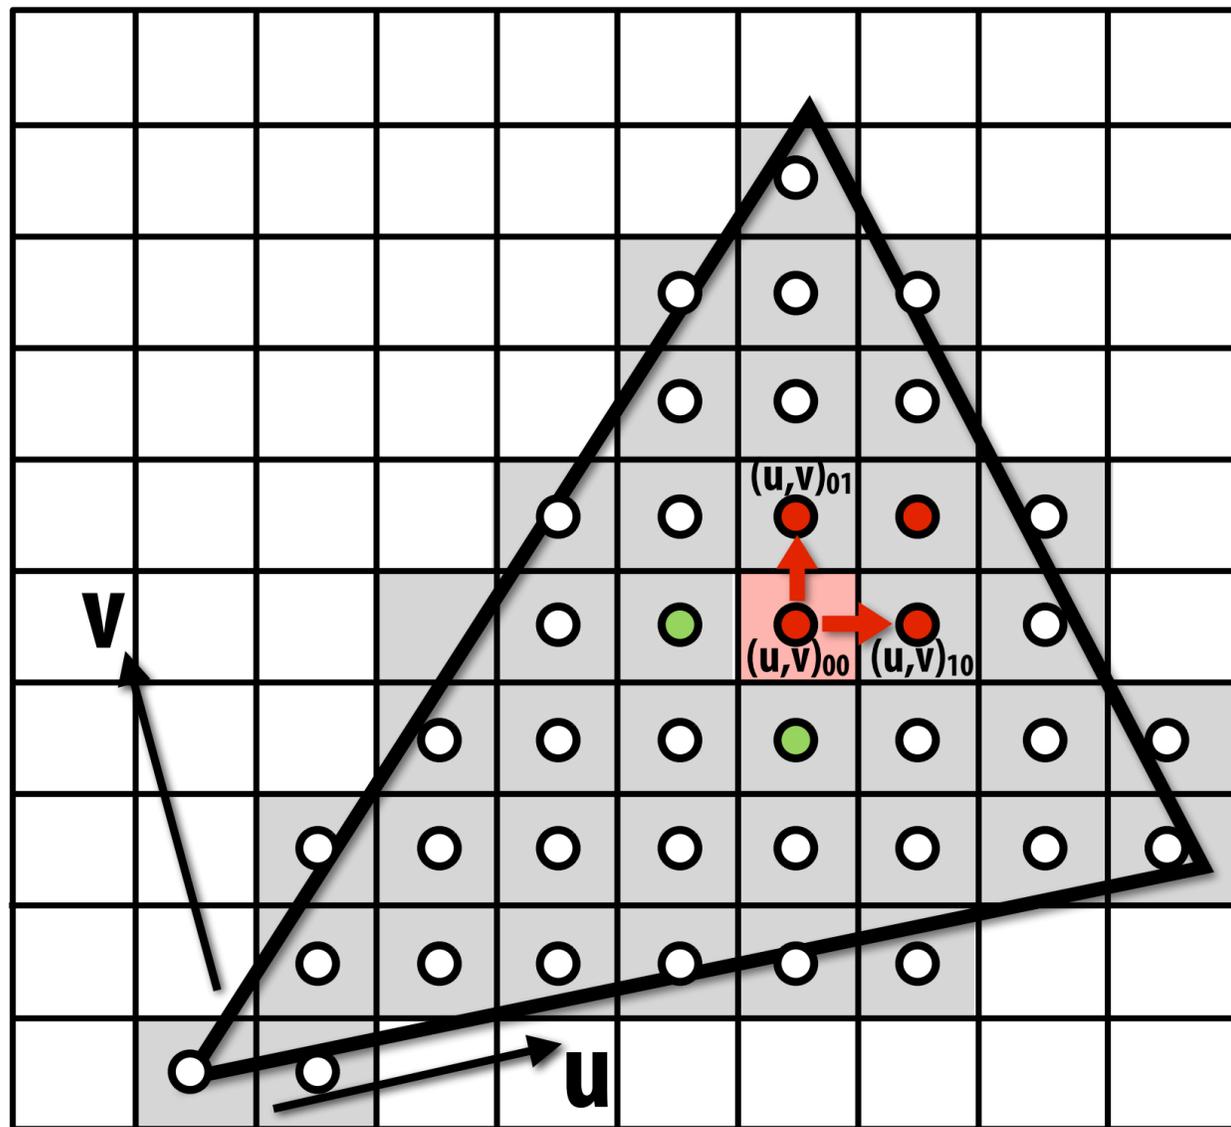
Screen space



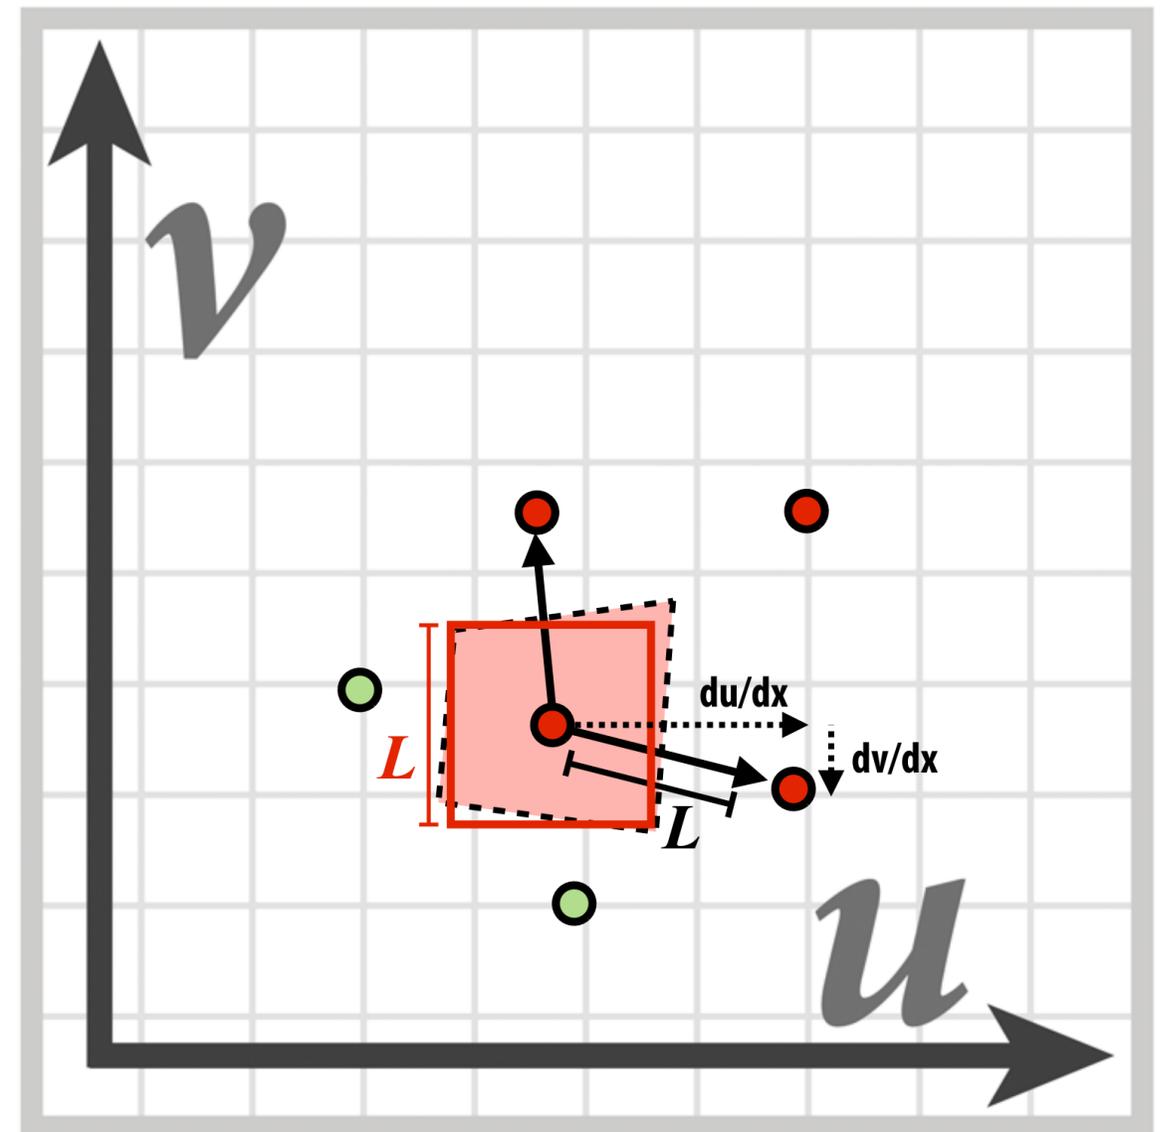
Texture space

Computing mipmap level

Compute differences between texture coordinate values of neighboring screen samples



$$\begin{aligned} \frac{du}{dx} &= u_{10} - u_{00} & \frac{dv}{dx} &= v_{10} - v_{00} \\ \frac{du}{dy} &= u_{01} - u_{00} & \frac{dv}{dy} &= v_{01} - v_{00} \end{aligned}$$



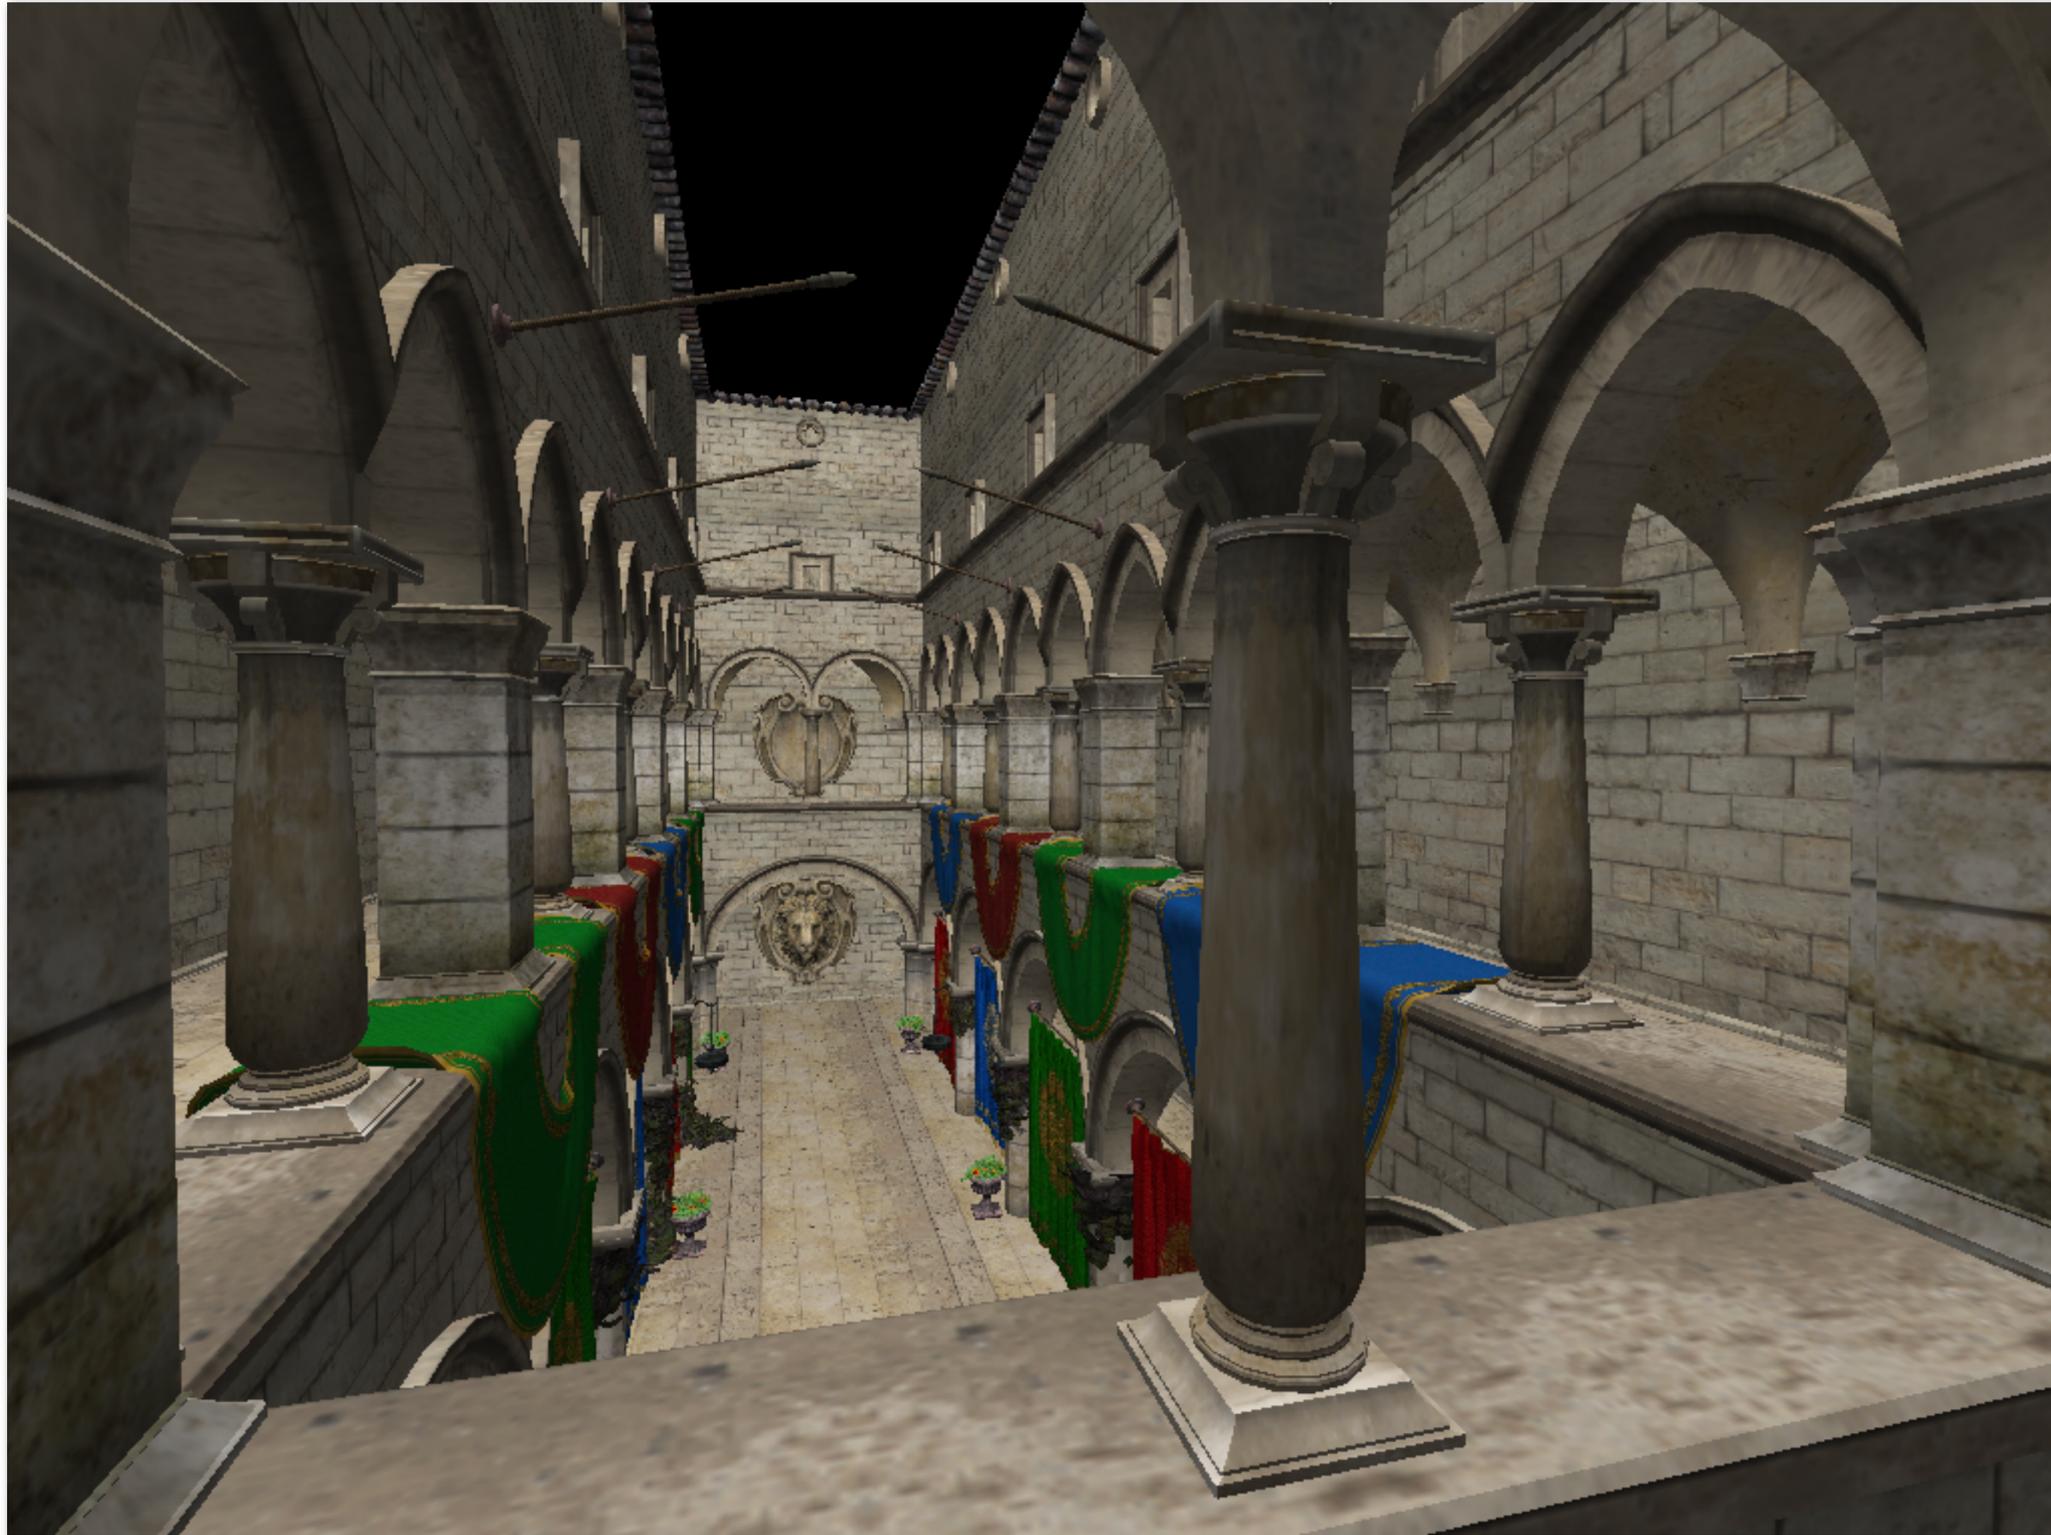
$$L = \max \left(\sqrt{\left(\frac{du}{dx}\right)^2 + \left(\frac{dv}{dx}\right)^2}, \sqrt{\left(\frac{du}{dy}\right)^2 + \left(\frac{dv}{dy}\right)^2} \right)$$

mip-map $d = \log_2 L$

Bilinear resampling at level 0



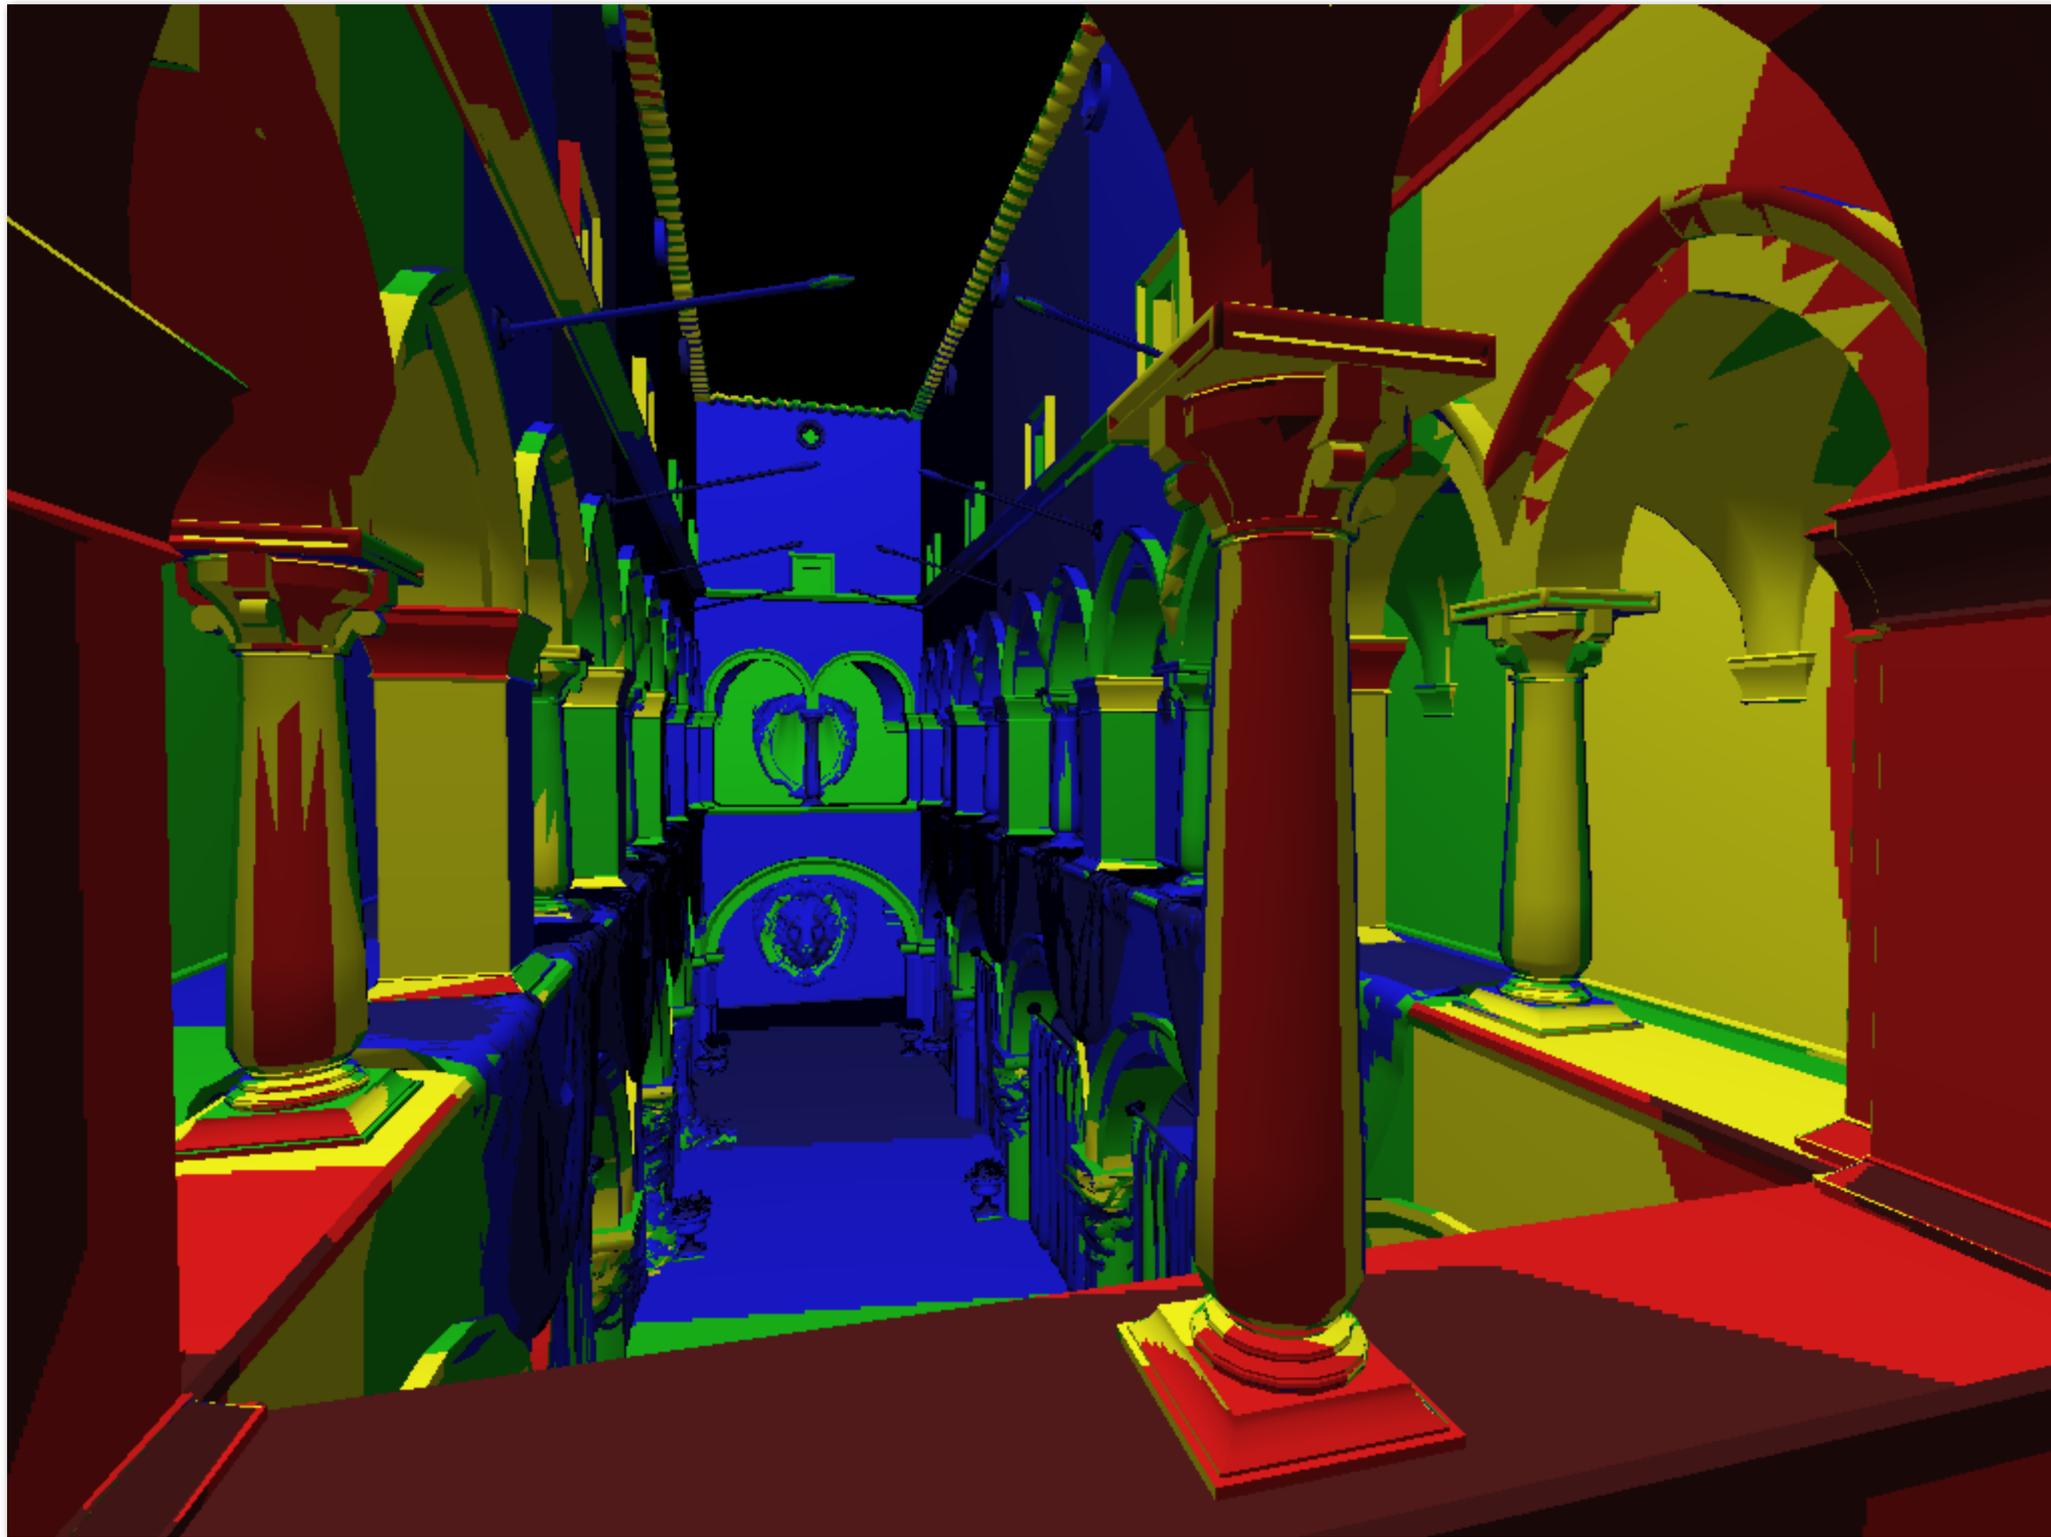
Bilinear resampling at level 2



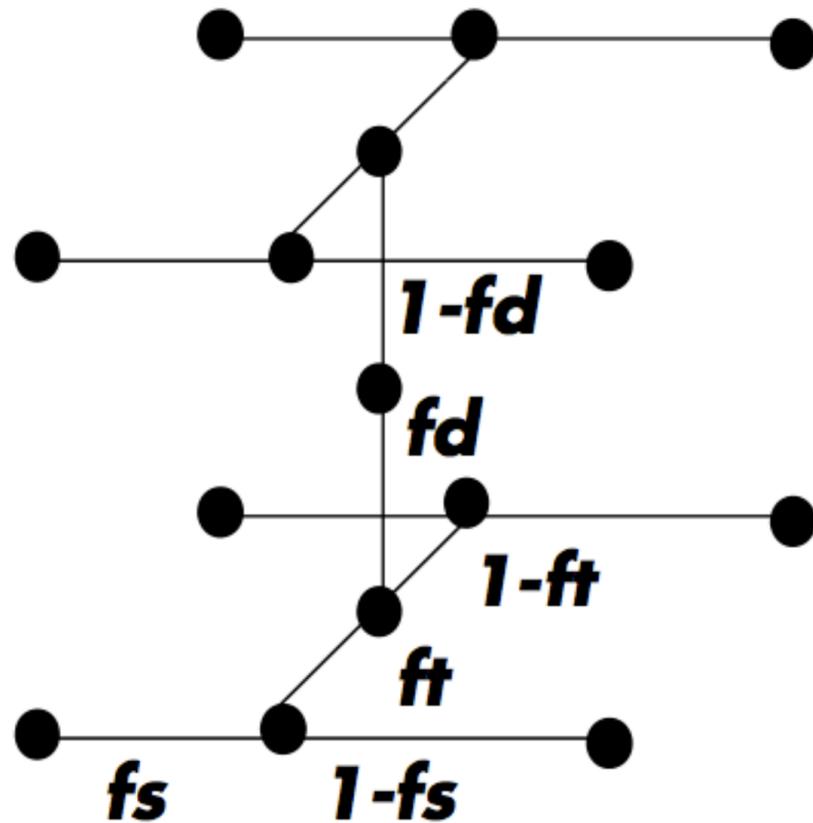
Bilinear resampling at level 4



Visualization of mipmap level (bilinear filtering only: d clamped to nearest level)



“Tri-linear” filtering



$$\text{lerp}(t, v_1, v_2) = v_1 + t(v_2 - v_1)$$

Bilinear resampling:

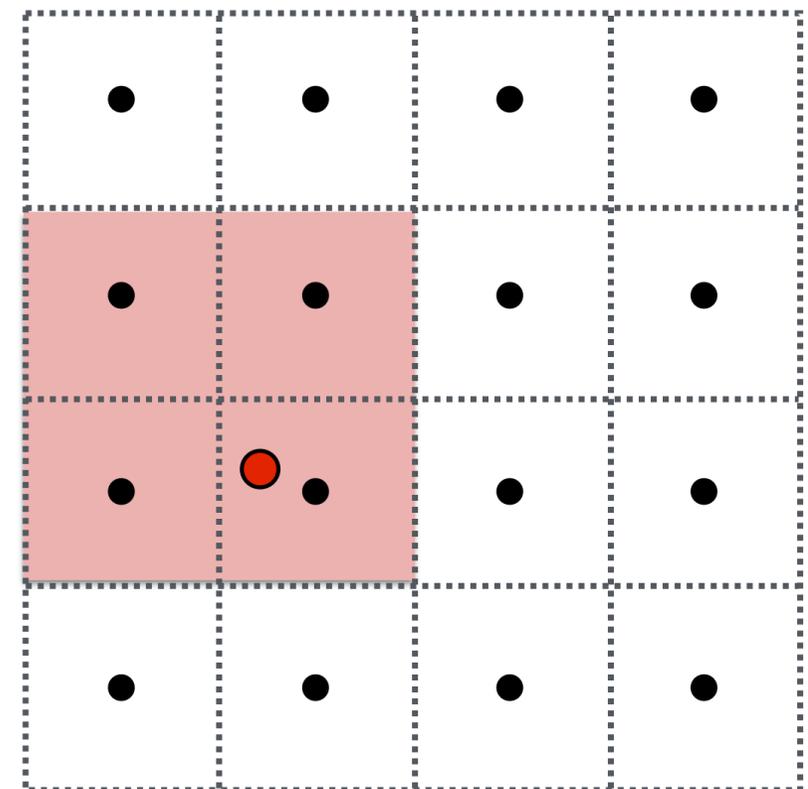
four texel reads

3 lerps (3 mul + 6 add)

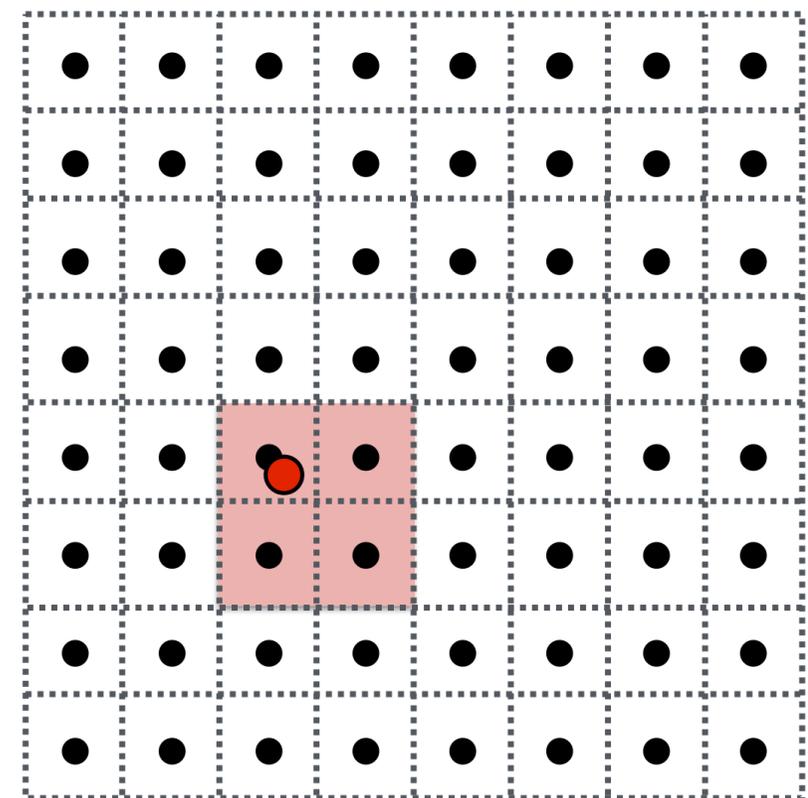
Trilinear resampling:

eight texel reads

7 lerps (7 mul + 14 add)

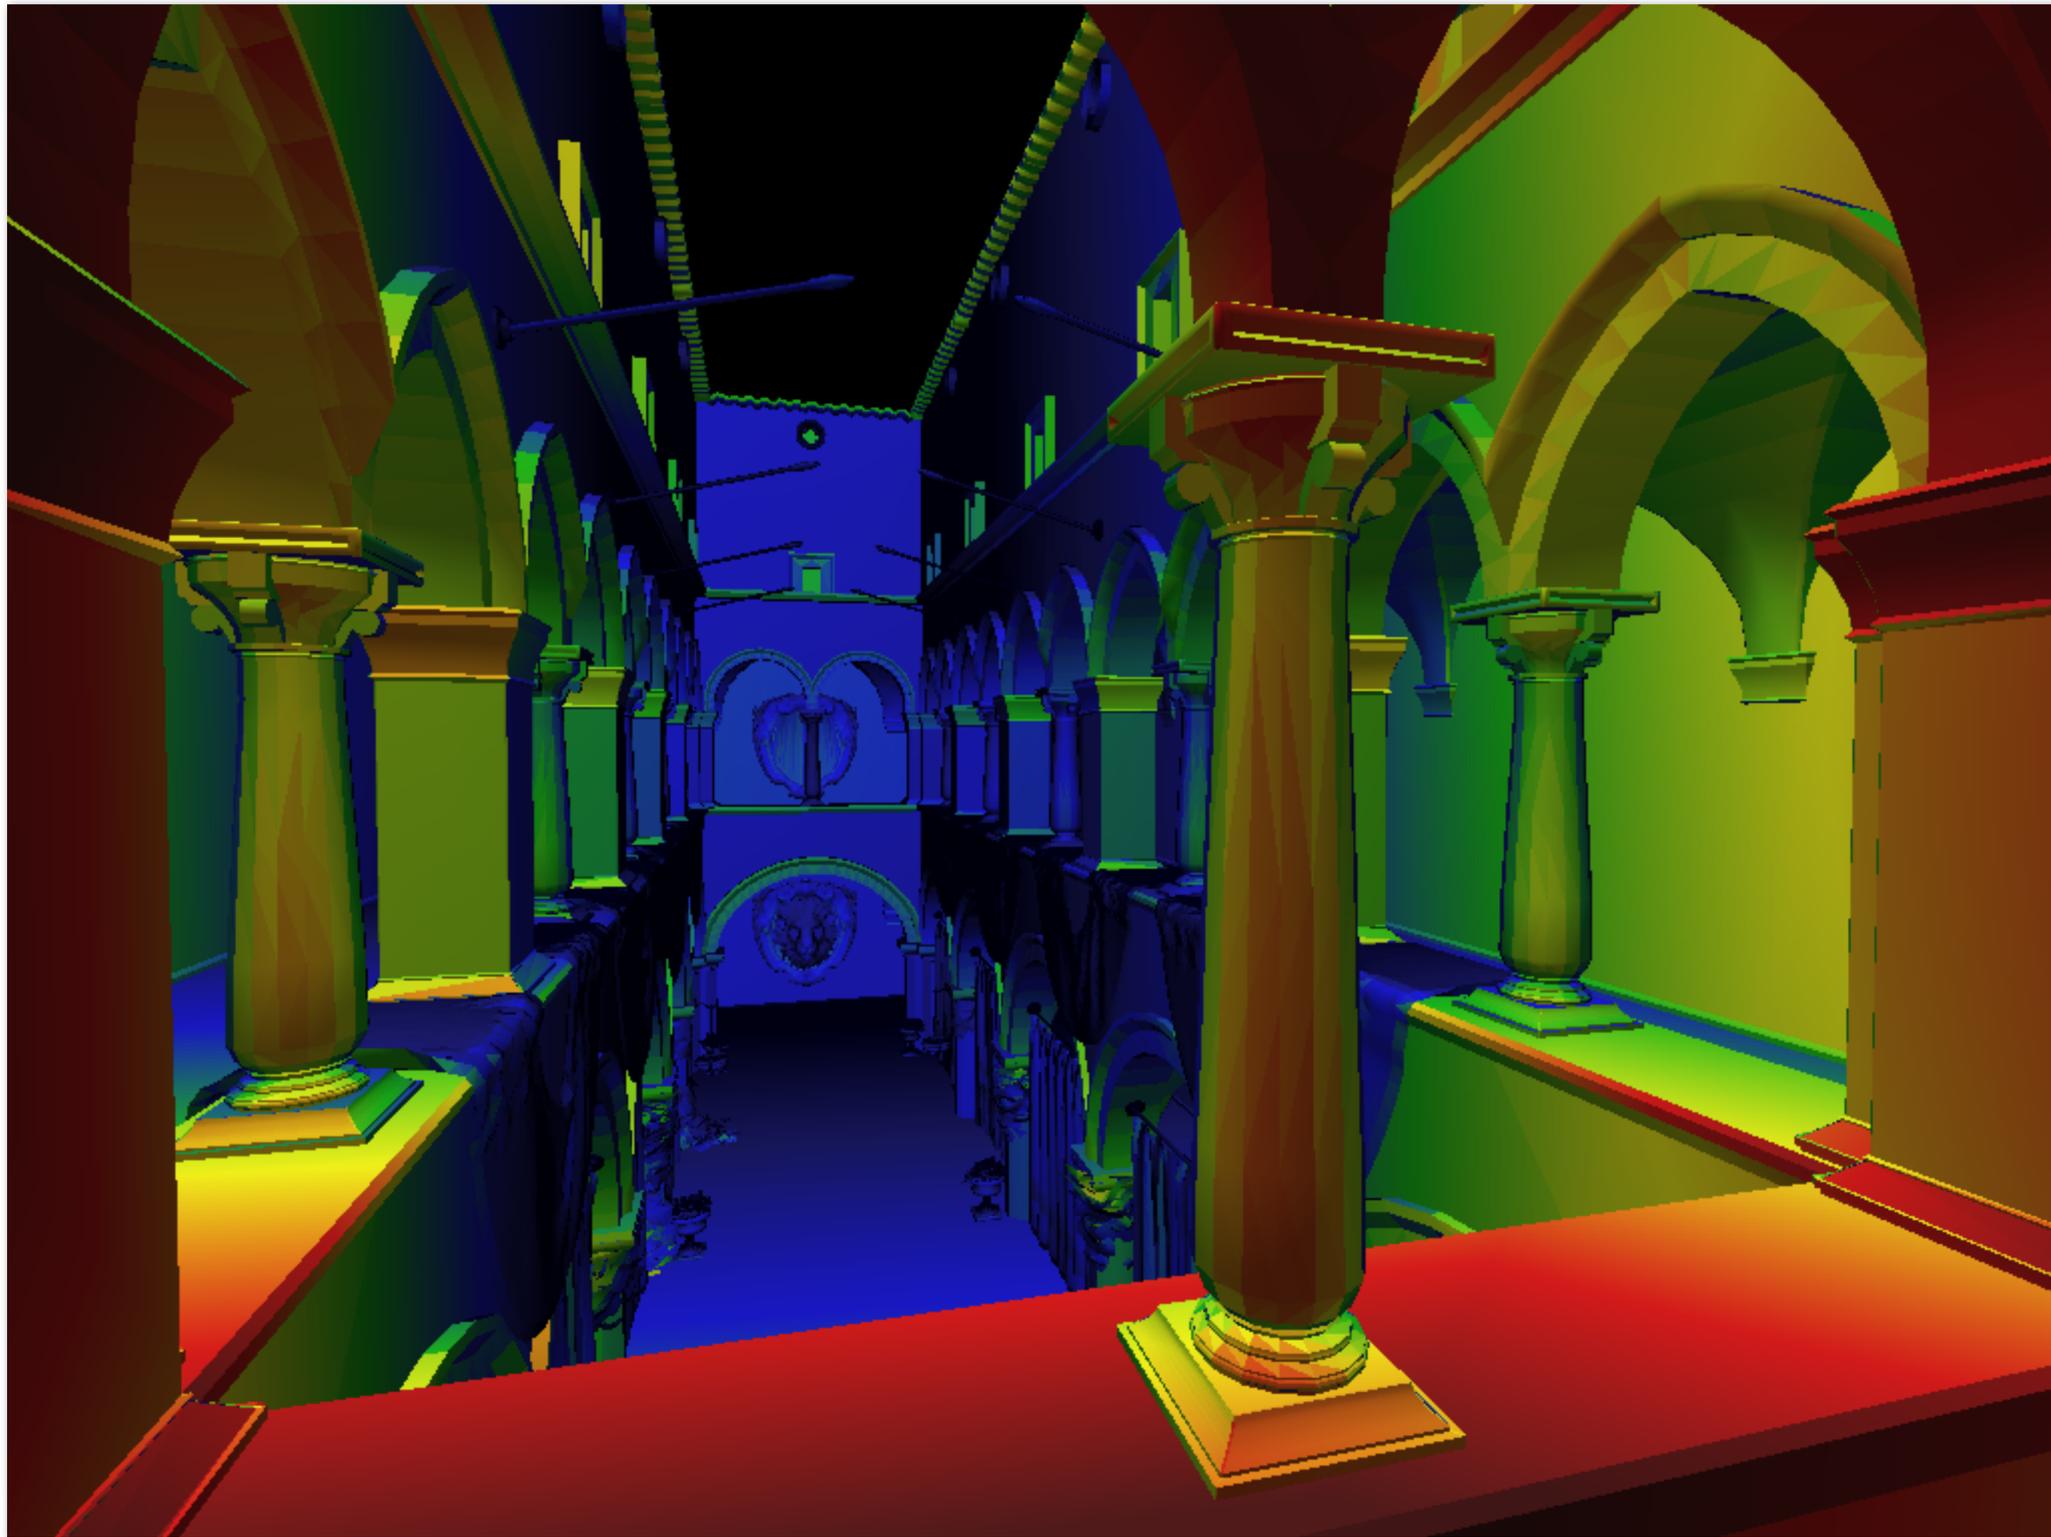


mip-map texels: level d+1



mip-map texels: level d

Visualization of mipmap level (trilinear filtering: visualization of continuous d)



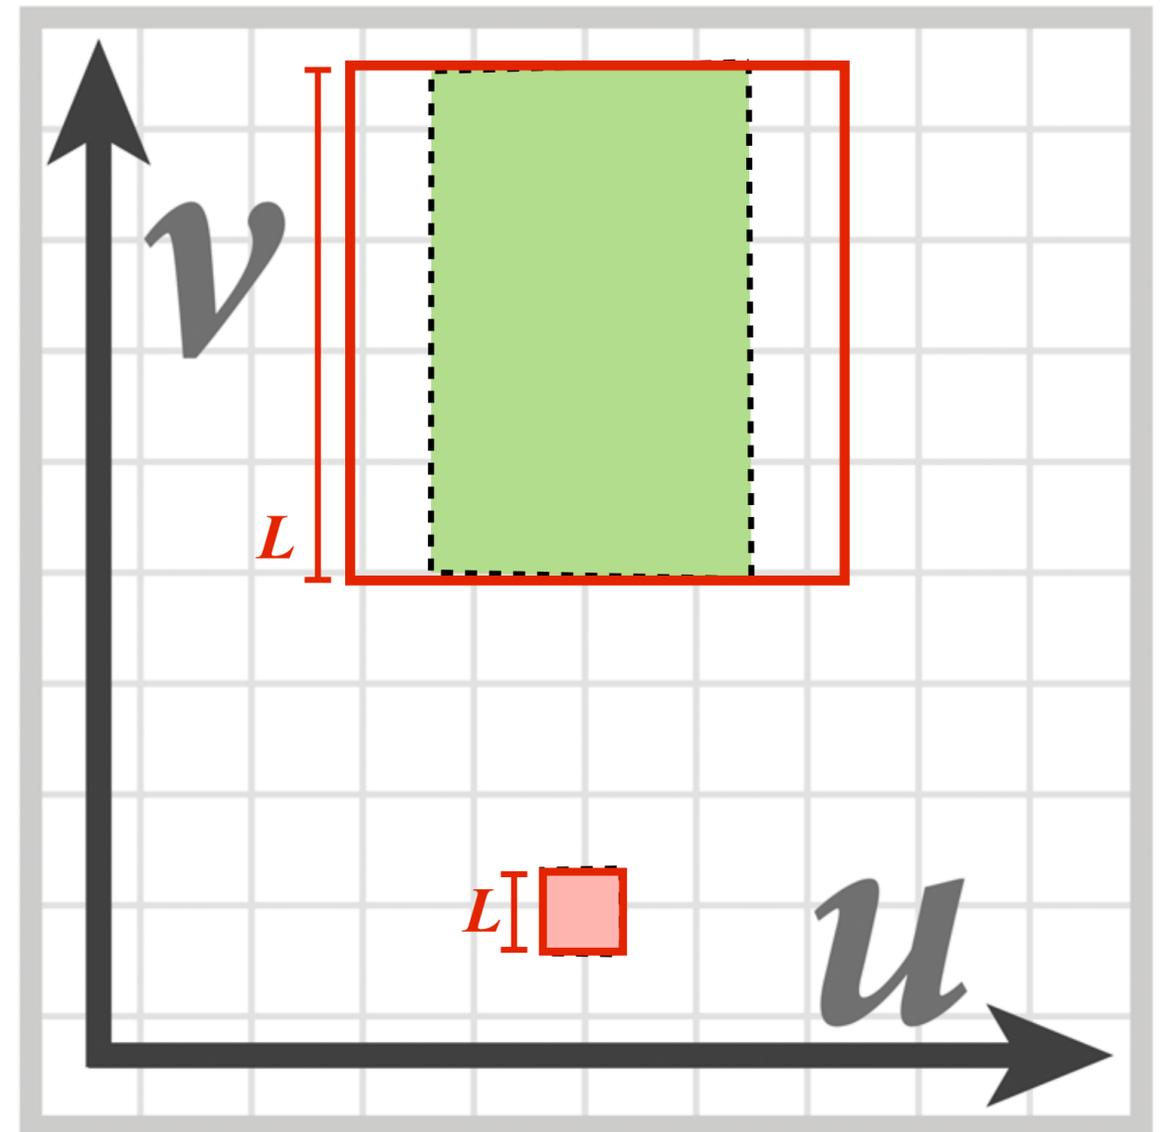
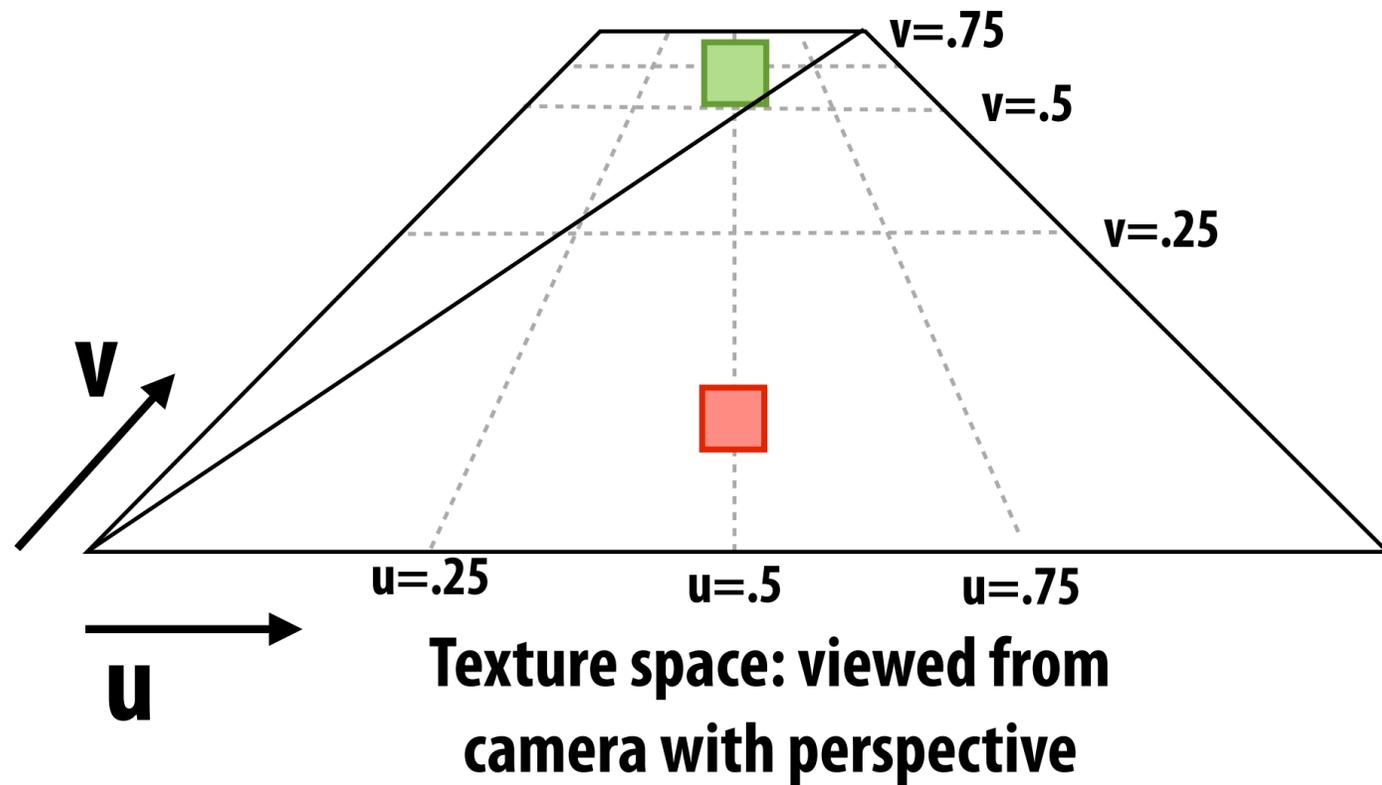
Bilinear vs trilinear filtering cost

- **Bilinear resampling:**
 - **4 texel reads**
 - **3 lerps (3 mul + 6 add)**

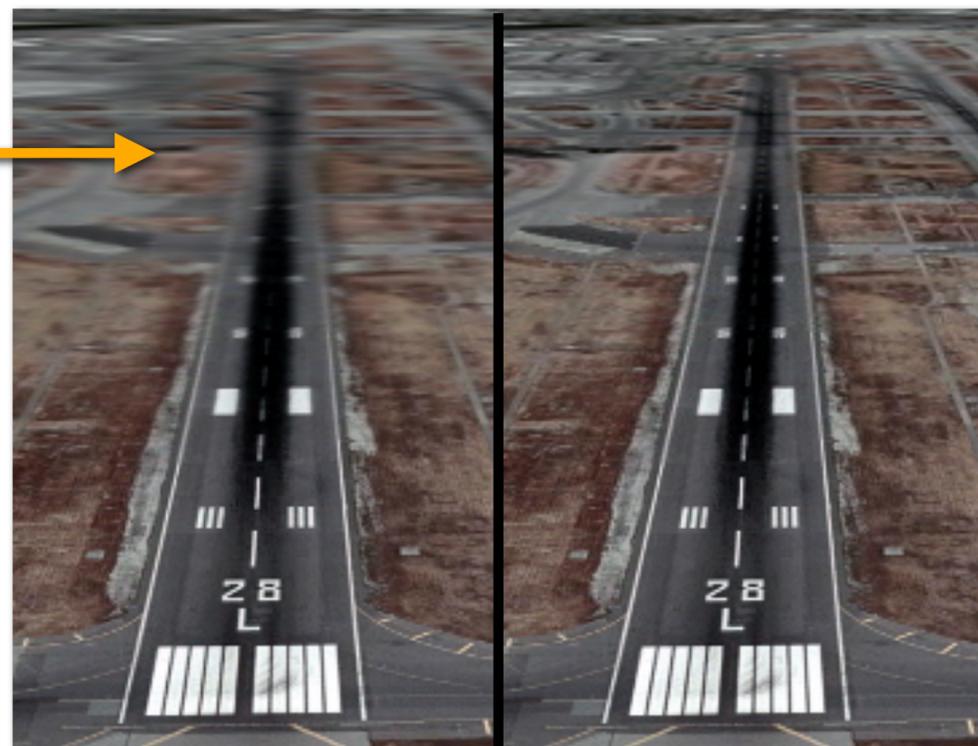
- **Trilinear resampling:**
 - **8 texel reads**
 - **7 lerps (7 mul + 14 add)**

Pixel area may not map to isotropic region in texture space

Proper filtering requires anisotropic filter footprint



Overblurring in u direction

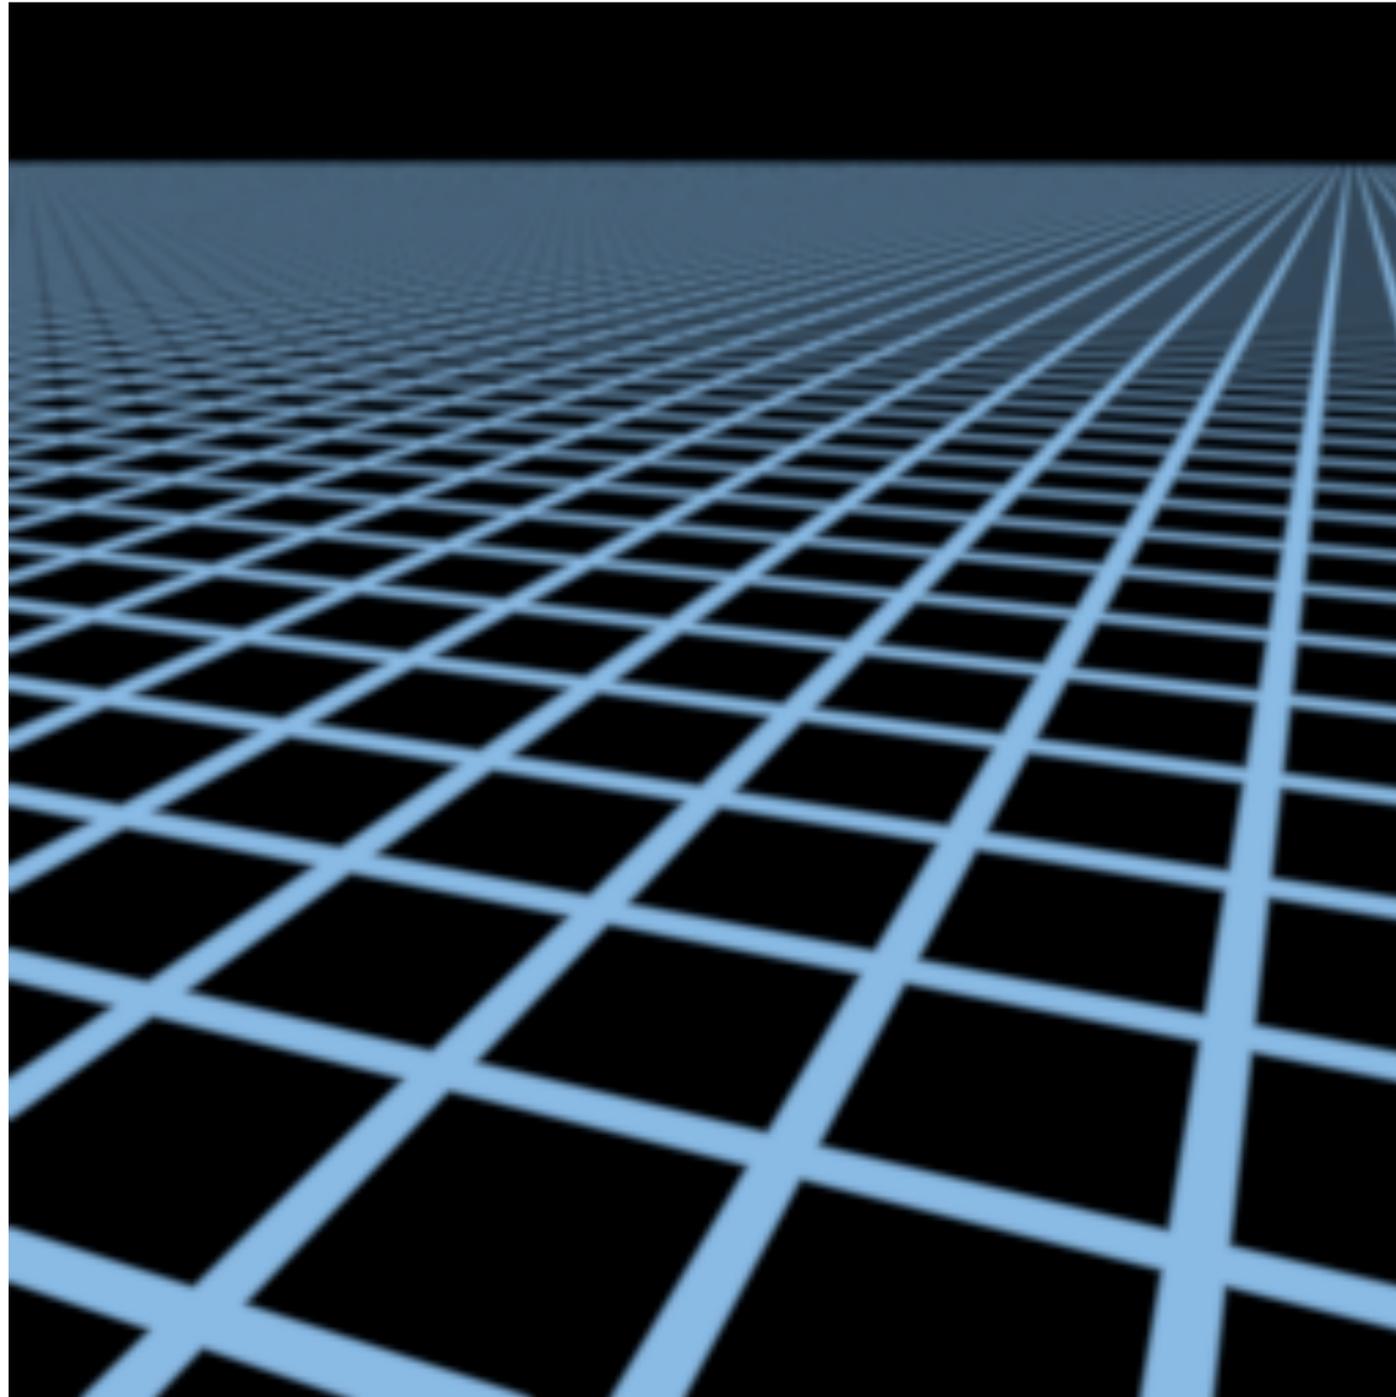


Trilinear (Isotropic) Filtering

Anisotropic Filtering

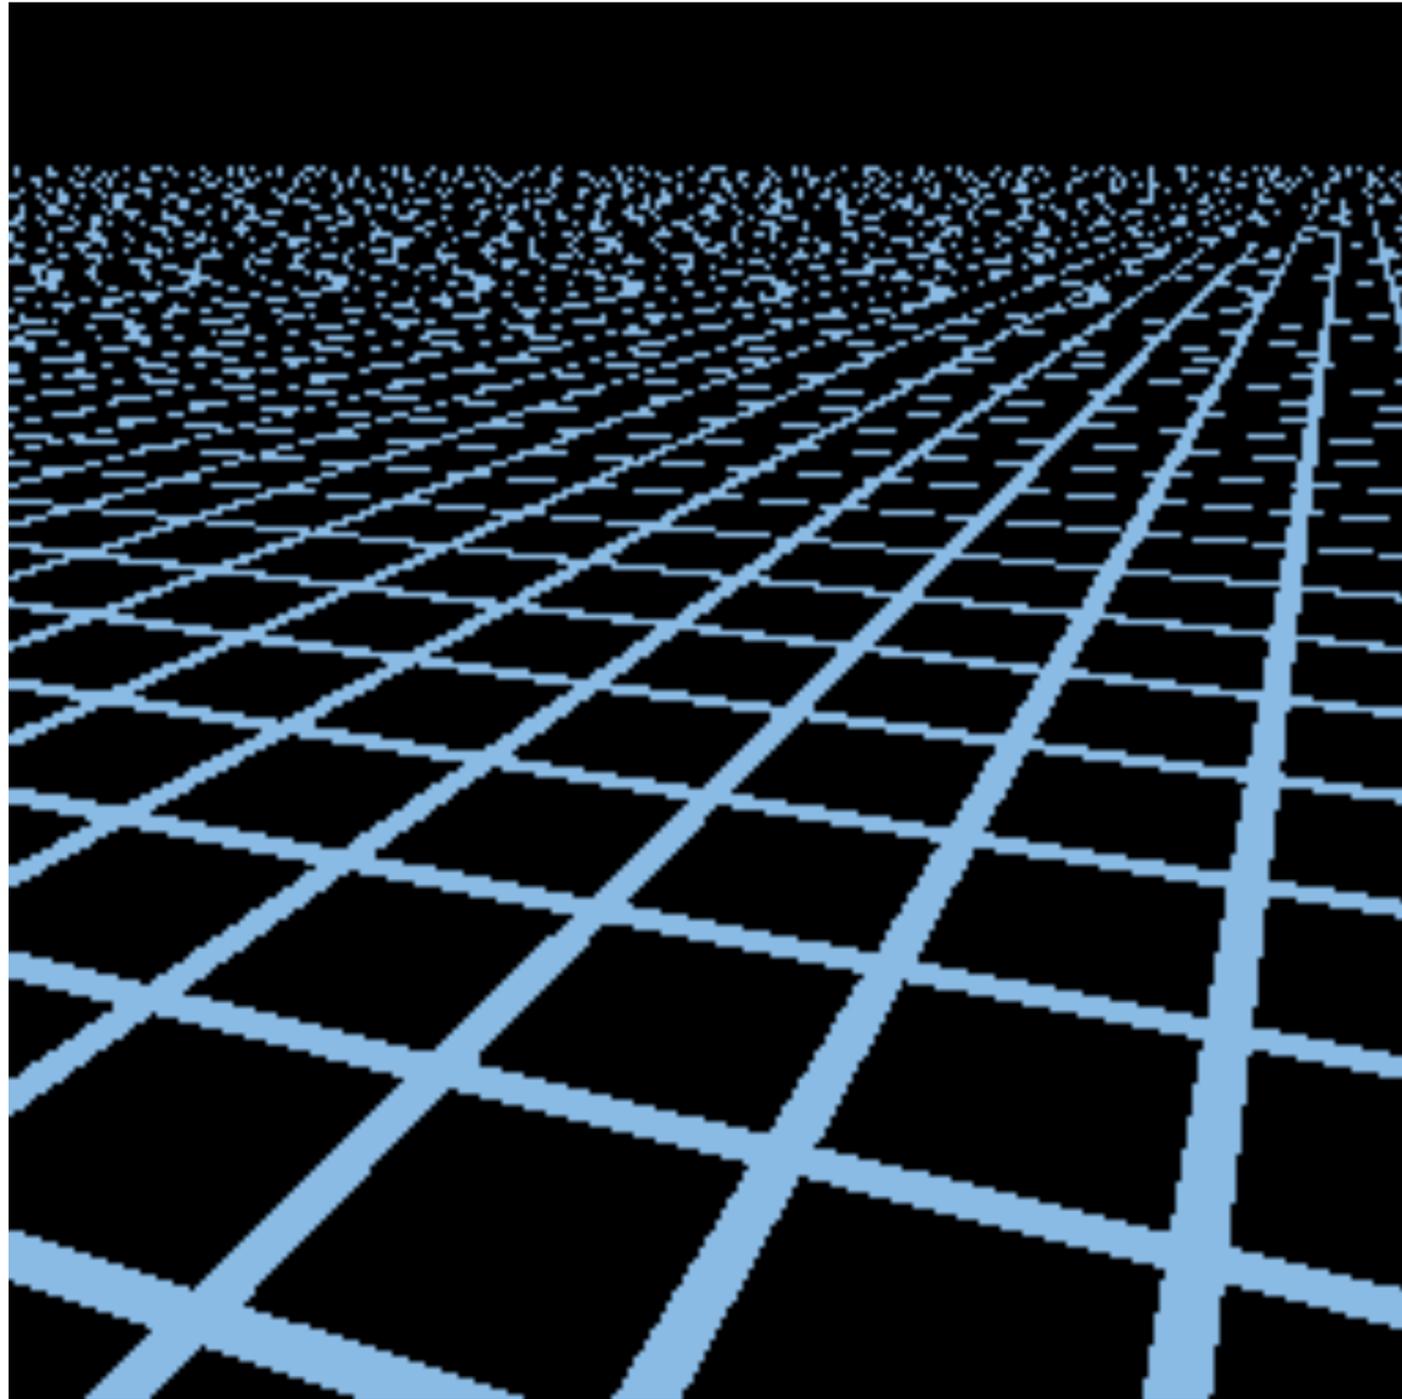
(Modern solutions combine multiple mip map samples)

Example: mipmap limitations



**Supersampling 512x
(desired answer)**

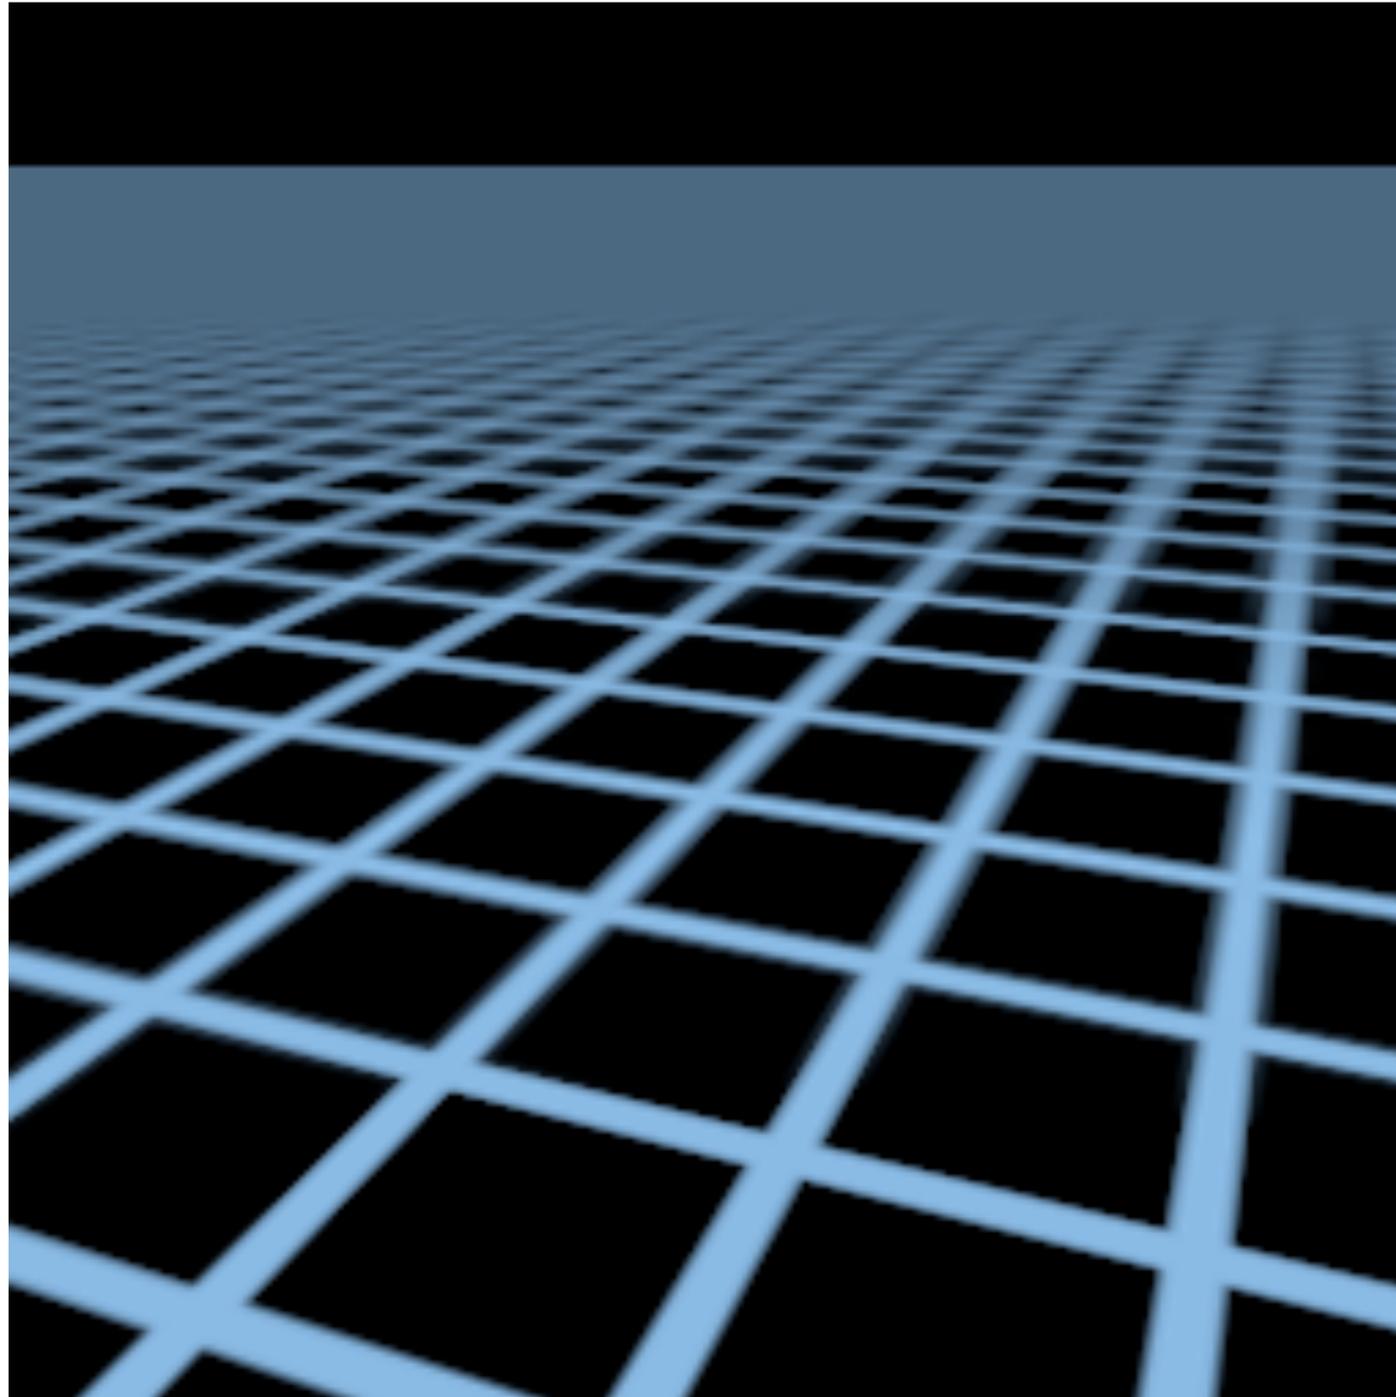
Example: mipmap limitations



Point sampling

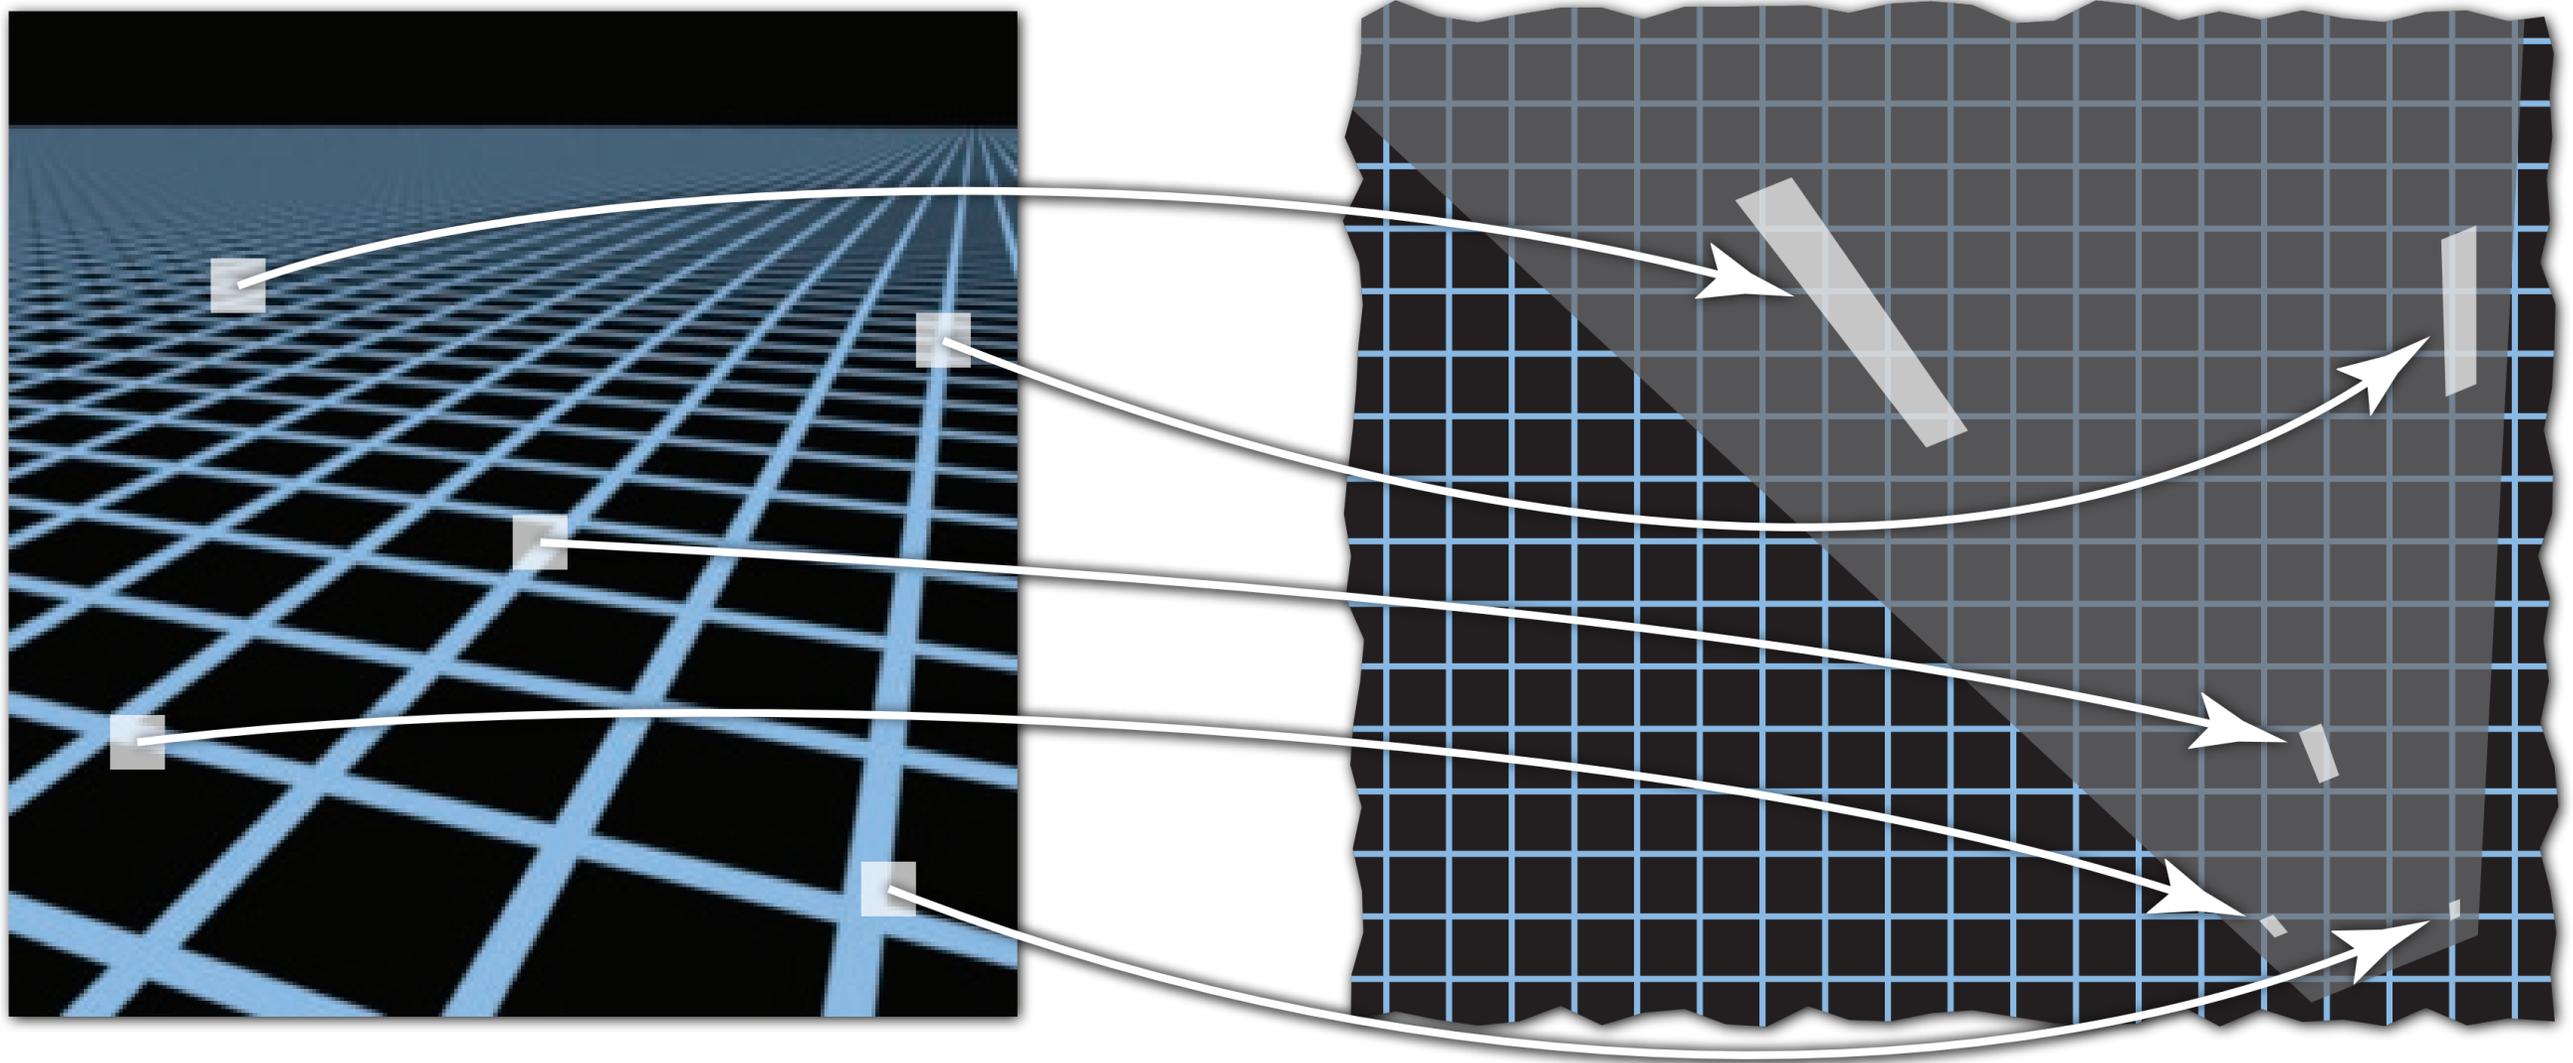
Example: mipmap limitations

Overblur
Why?



Mipmap trilinear sampling

Screen pixel footprint in texture space

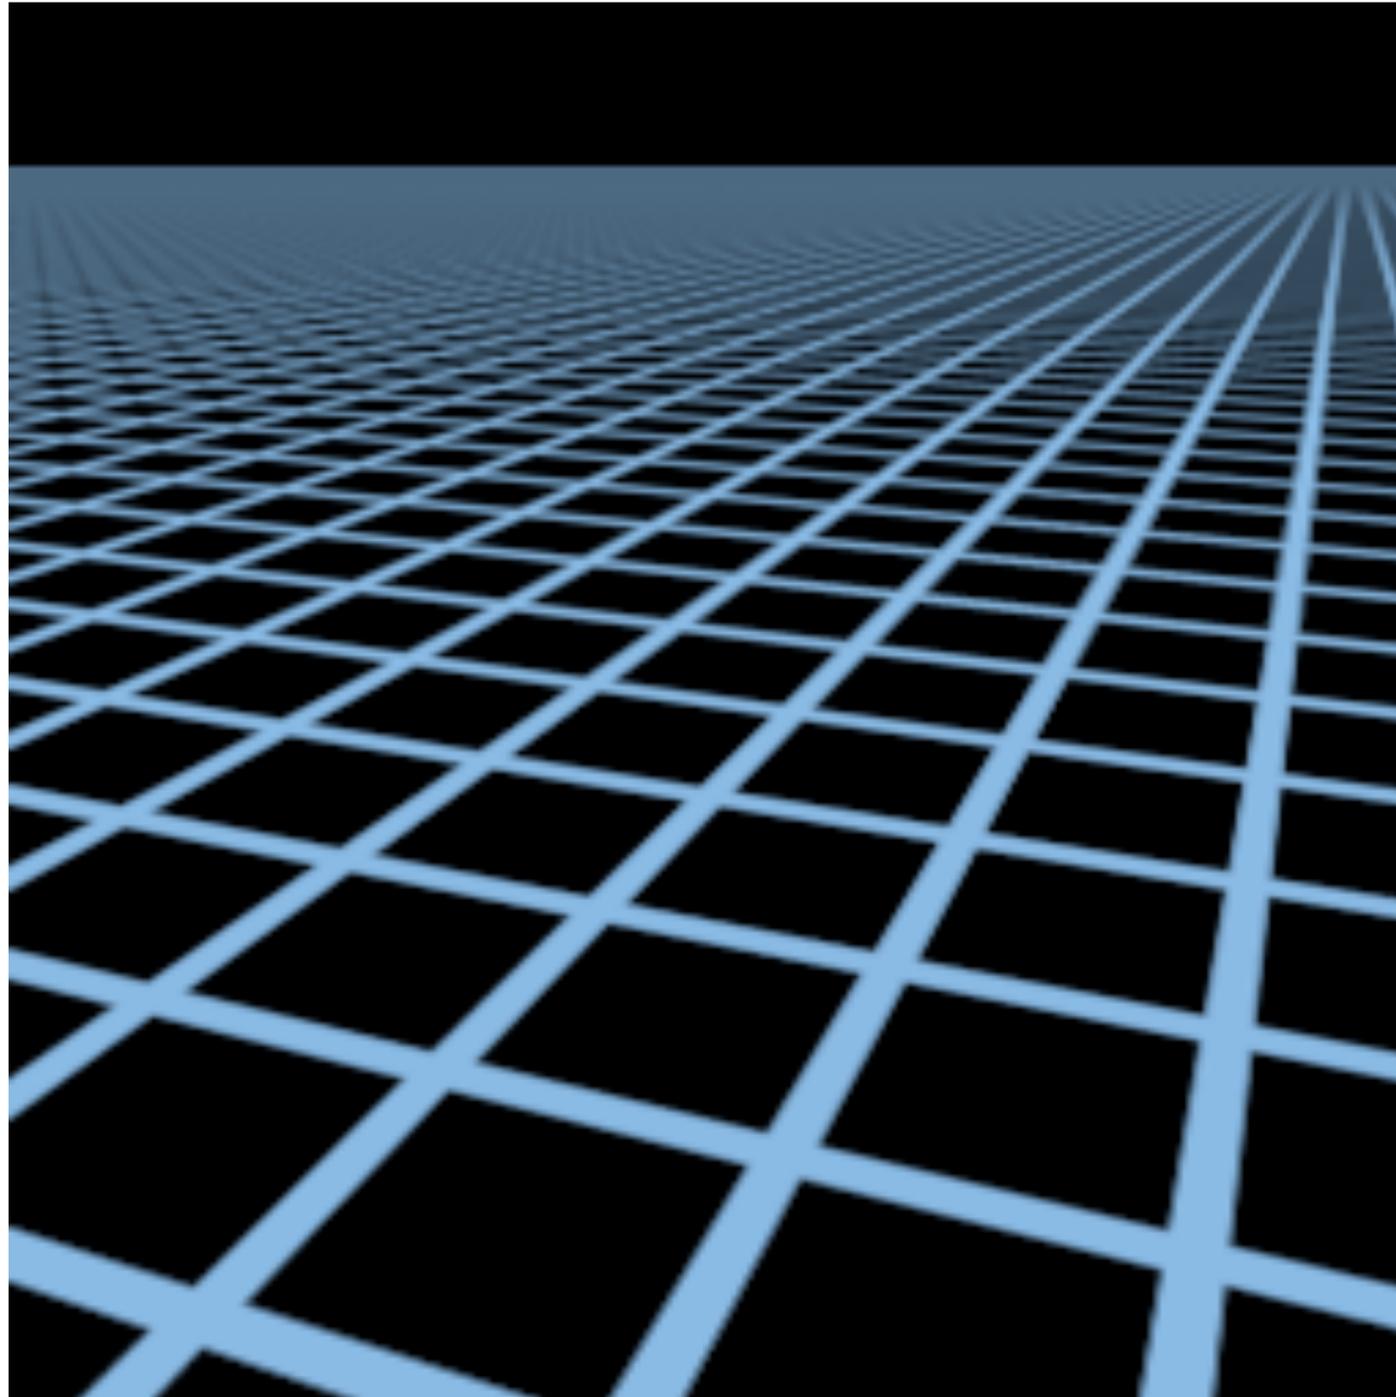


Screen space

Texture space

Texture sampling pattern not rectilinear or isotropic

Anisotropic filtering



Elliptical weighted average (EWA) filtering
(uses multiple lookups into mip-map to approximate filter region)

Summary: texture filtering using the mip map

- **Small storage overhead (33%)**
 - Mipmap is $4/3$ the size of original texture image
- **For each isotropically-filtered sampling operation**
 - Constant filtering cost (independent of mip map level)
 - Constant number of texels accessed (independent of mip map level)
- **Combat aliasing with *prefiltering*, rather than supersampling**
 - Recall: we used supersampling to address aliasing problem when sampling coverage
- **Bilinear/trilinear filtering is isotropic and thus will “overblur” to avoid aliasing**
 - Anisotropic texture filtering provides higher image quality at higher compute and memory bandwidth cost (in practice: multiple mip map samples)

A full texture sampling operation

1. Compute u and v from screen sample x,y (via evaluation of attribute equations)
2. Compute $du/dx, du/dy, dv/dx, dv/dy$ differentials from screen-adjacent samples.
3. Compute mip map level d
4. Convert normalized $[0,1]$ texture coordinate (u,v) to texture coordinates U,V in $[W,H]$
5. Compute required texels in window of filter
6. Load required texels from memory (need eight texels for trilinear)
7. Perform tri-linear interpolation according to (U, V, d)

Takeaway: a texture sampling operation is not just an image pixel lookup! It involves a significant amount of math.

For this reason, modern GPUs have dedicated fixed-function hardware support for performing texture sampling operations.

Texturing summary

- **Texture coordinates: define mapping between points on triangle's surface (object coordinate space) to points in texture coordinate space**
- **Texture mapping is a sampling operation and is prone to aliasing**
 - **Solution: prefilter texture map to eliminate high frequencies in texture signal**
 - **Mipmap: precompute and store multiple multiple resampled versions of the texture image (each with different amounts of low-pass filtering)**
 - **During rendering: dynamically select how much low-pass filtering is required based on distance between neighboring screen samples in texture space**
 - **Goal is to retain as much high-frequency content (detail) in the texture as possible, while avoiding aliasing**

Acknowledgements

- **Thanks to Ren Ng, Pat Hanrahan, and Keenan Crane for slide materials**