

Lecture 6: Texture Mapping

**Computer Graphics: Rendering,
Geometry, and Image Manipulation
Stanford CS248A, Winter 2026**

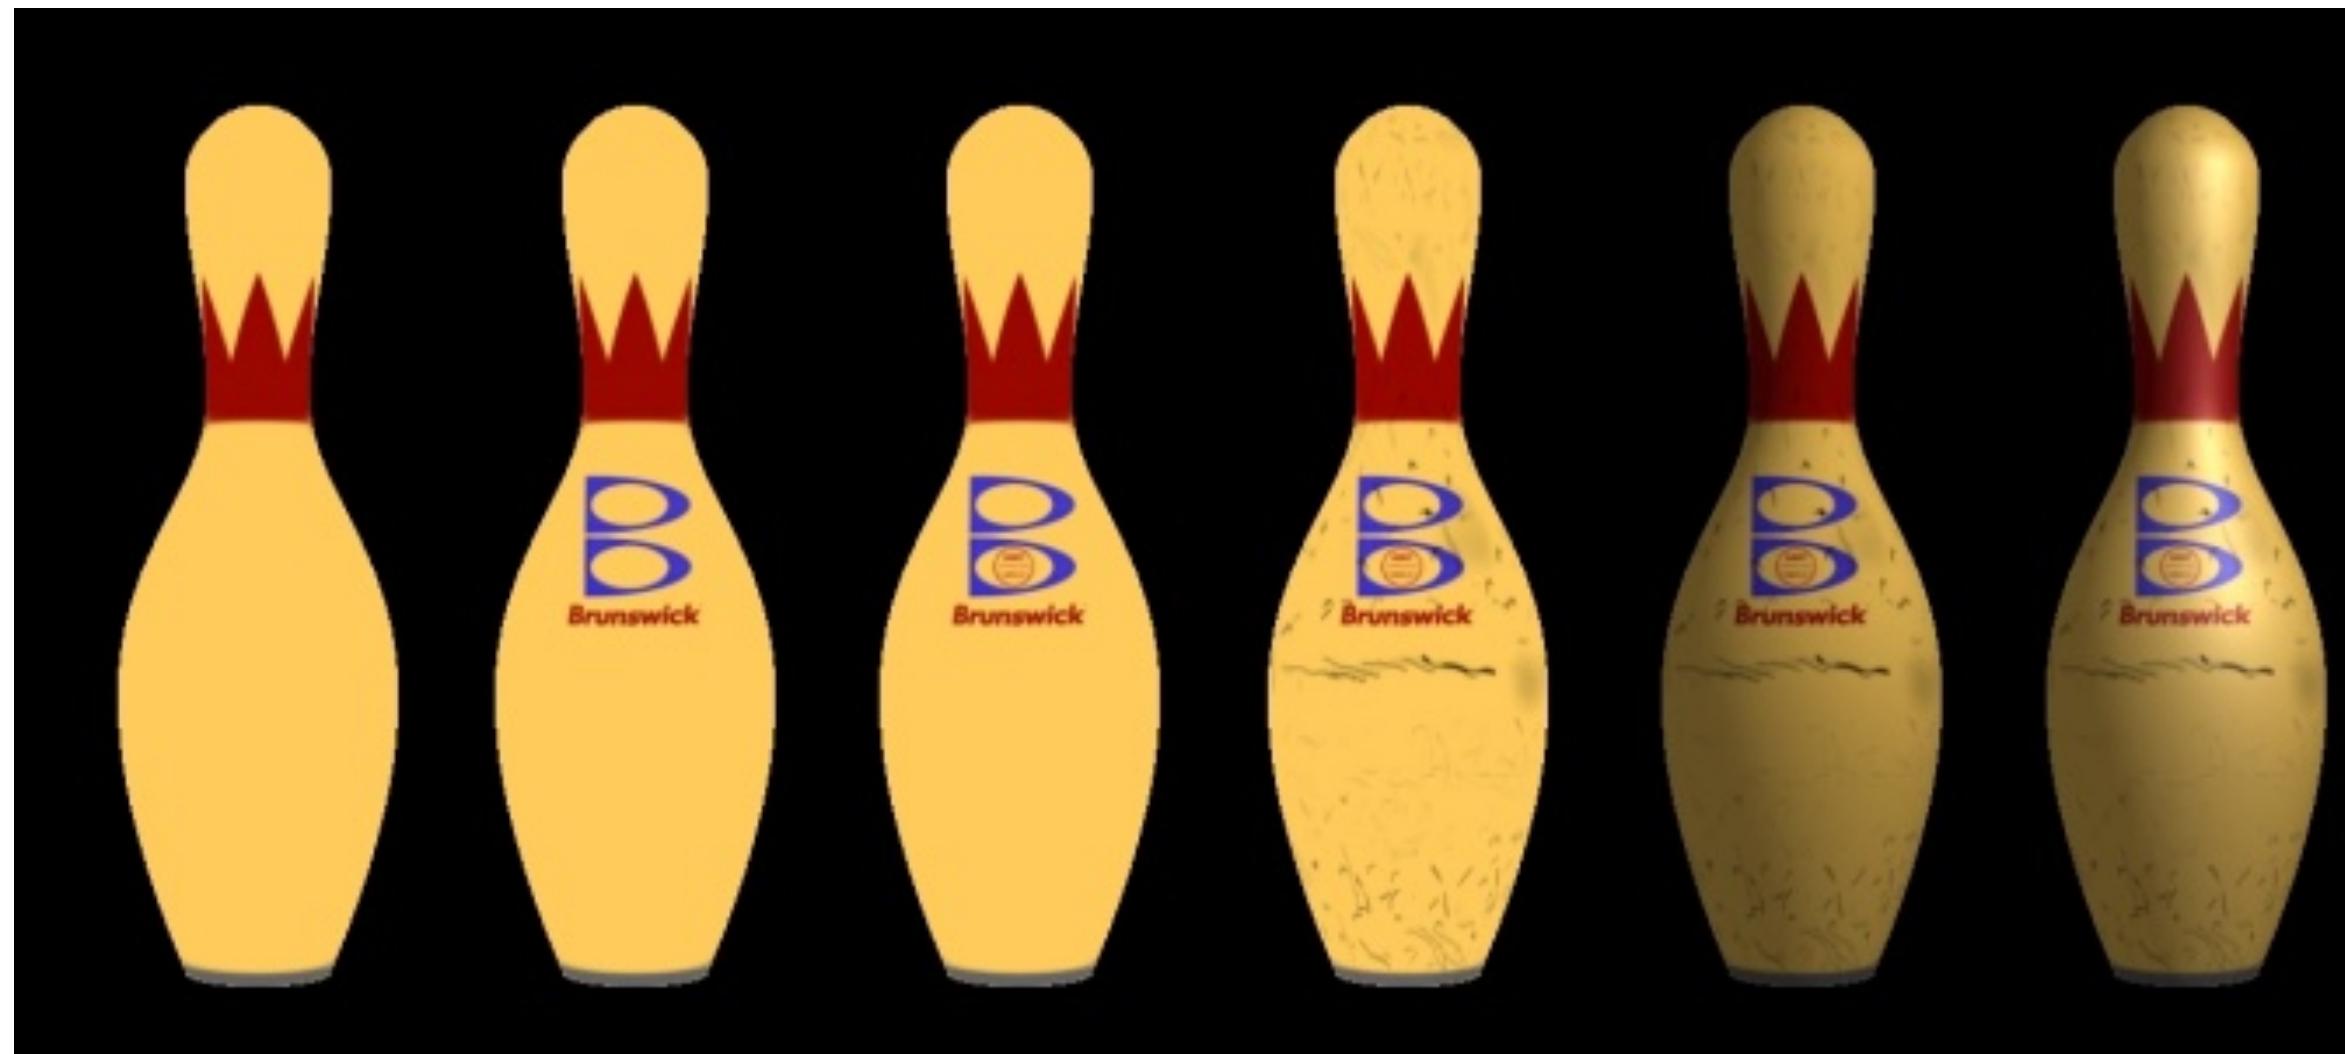


Many uses of texture mapping

Define variation in surface reflectance



Describe surface material properties



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

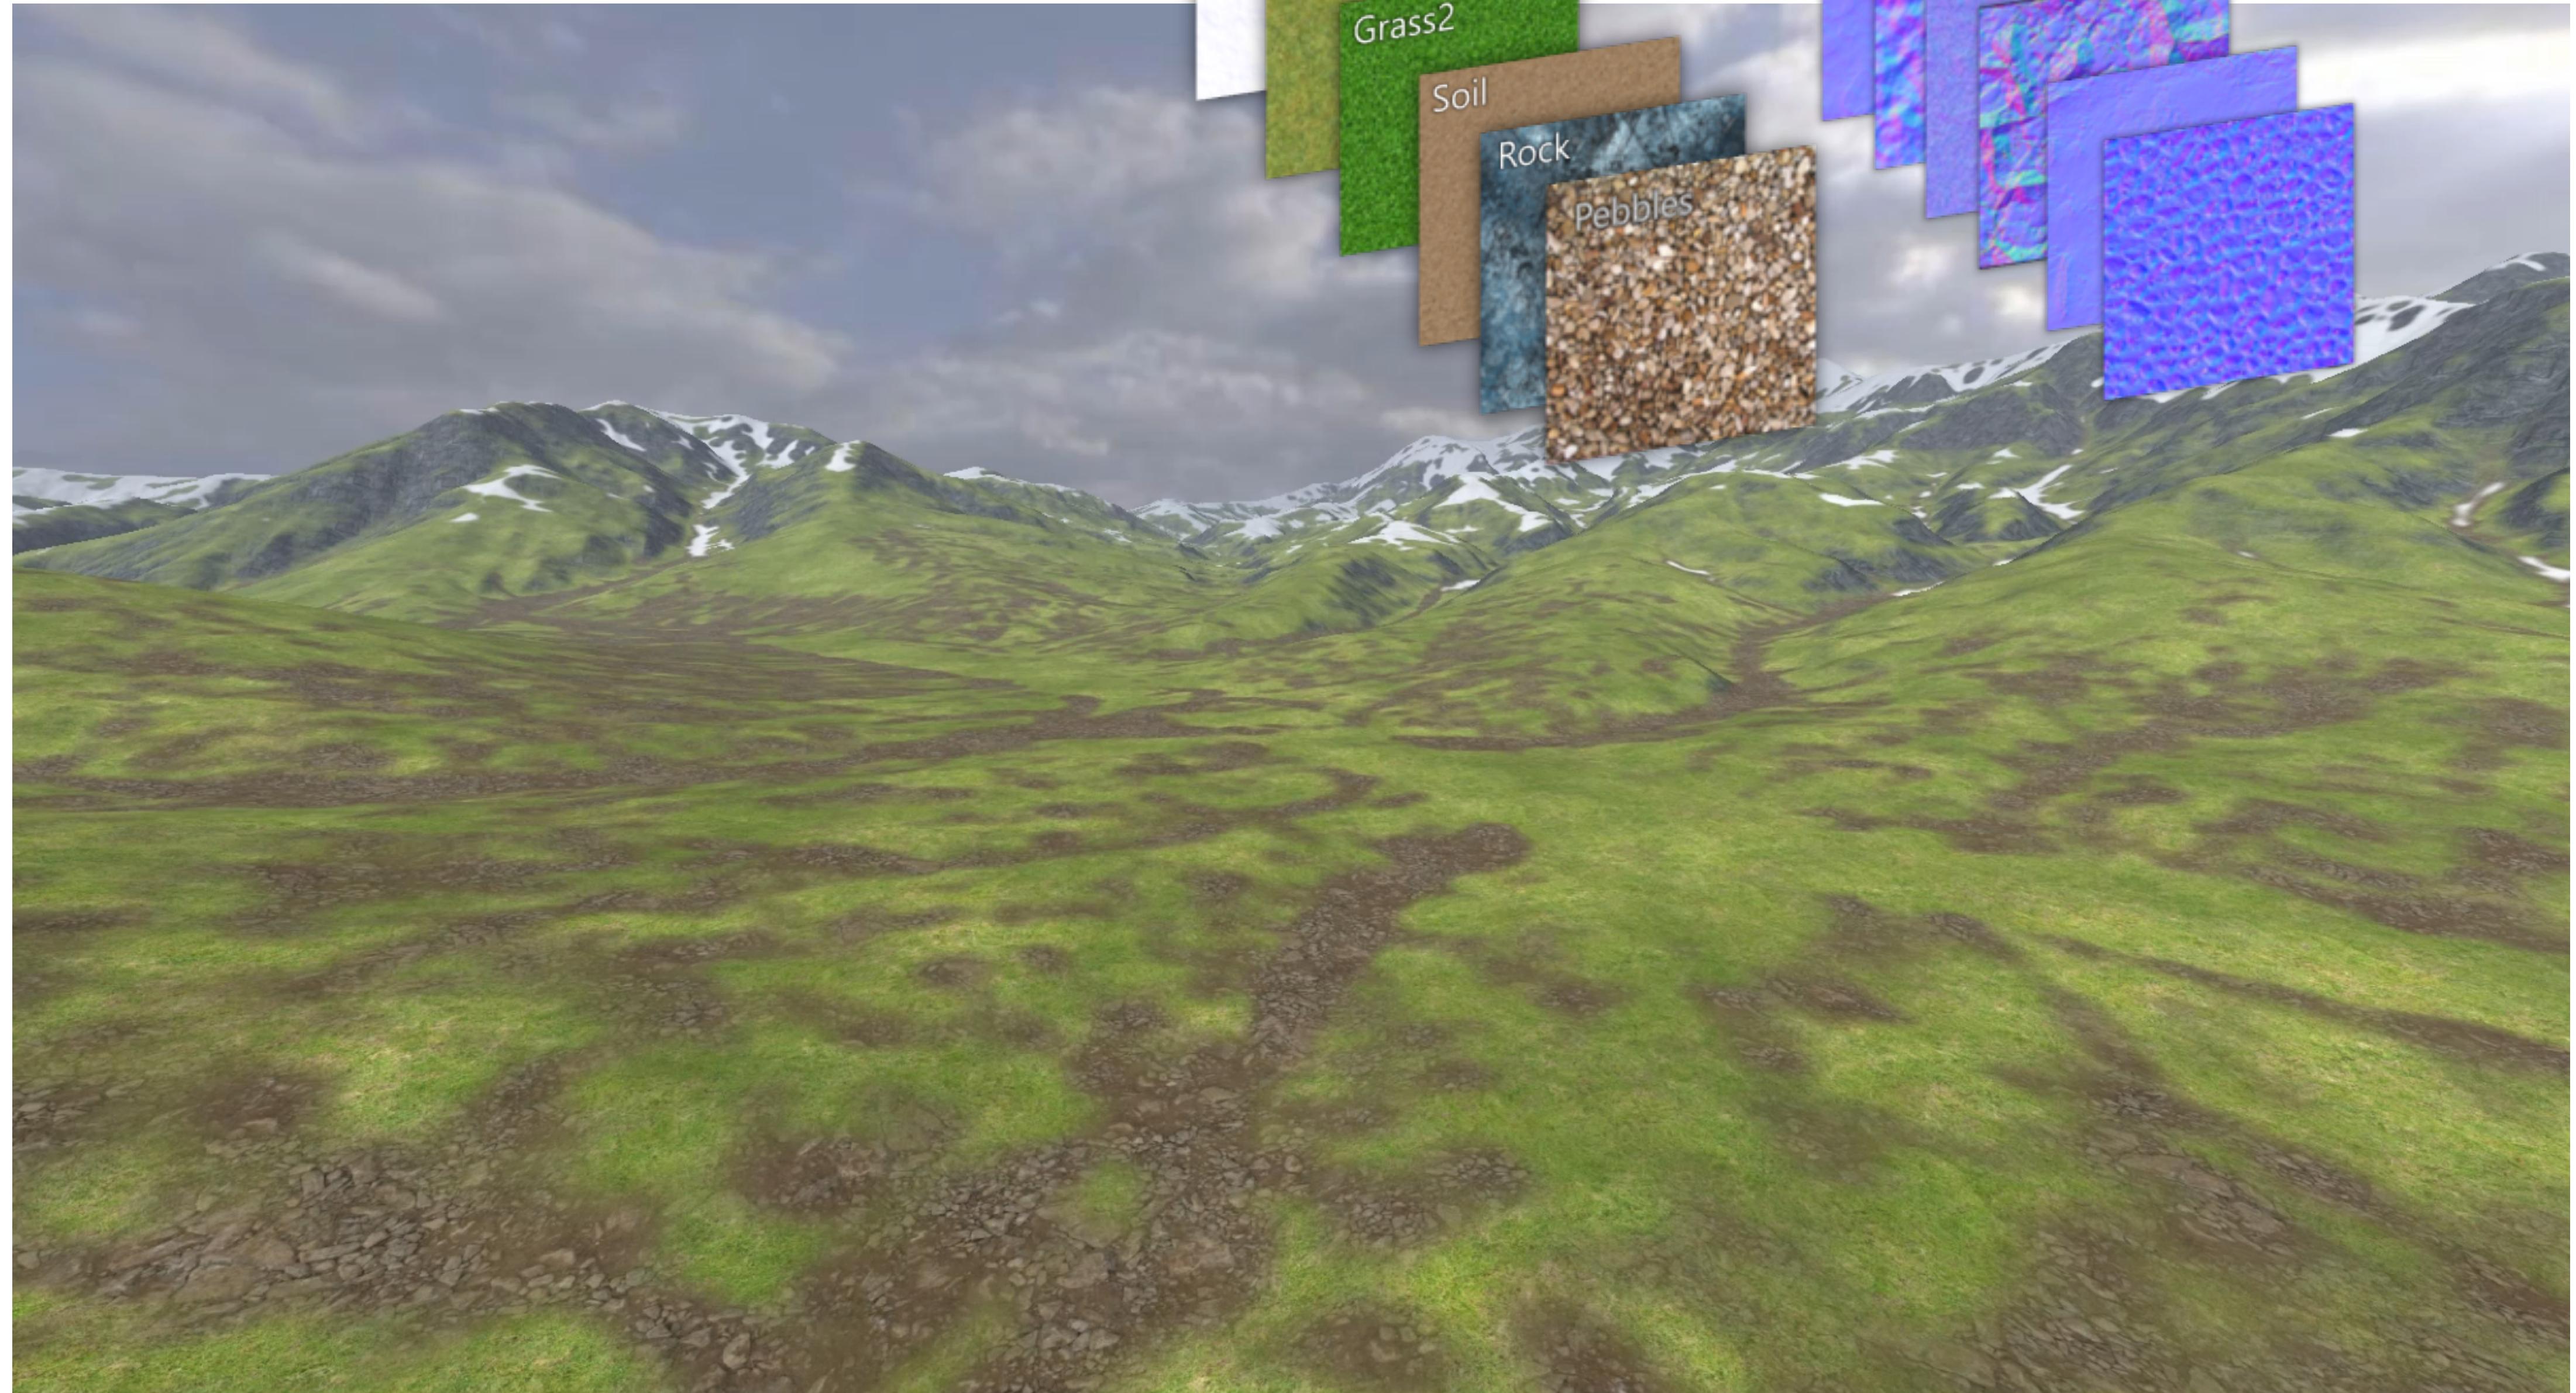


(C)2013 CRYTEK GMBH. ALL RIGHTS RESERVED. RYSE IS A REGISTERED TRADEMARK OF CRYTEK GMBH

RYSE
SON OF ROME

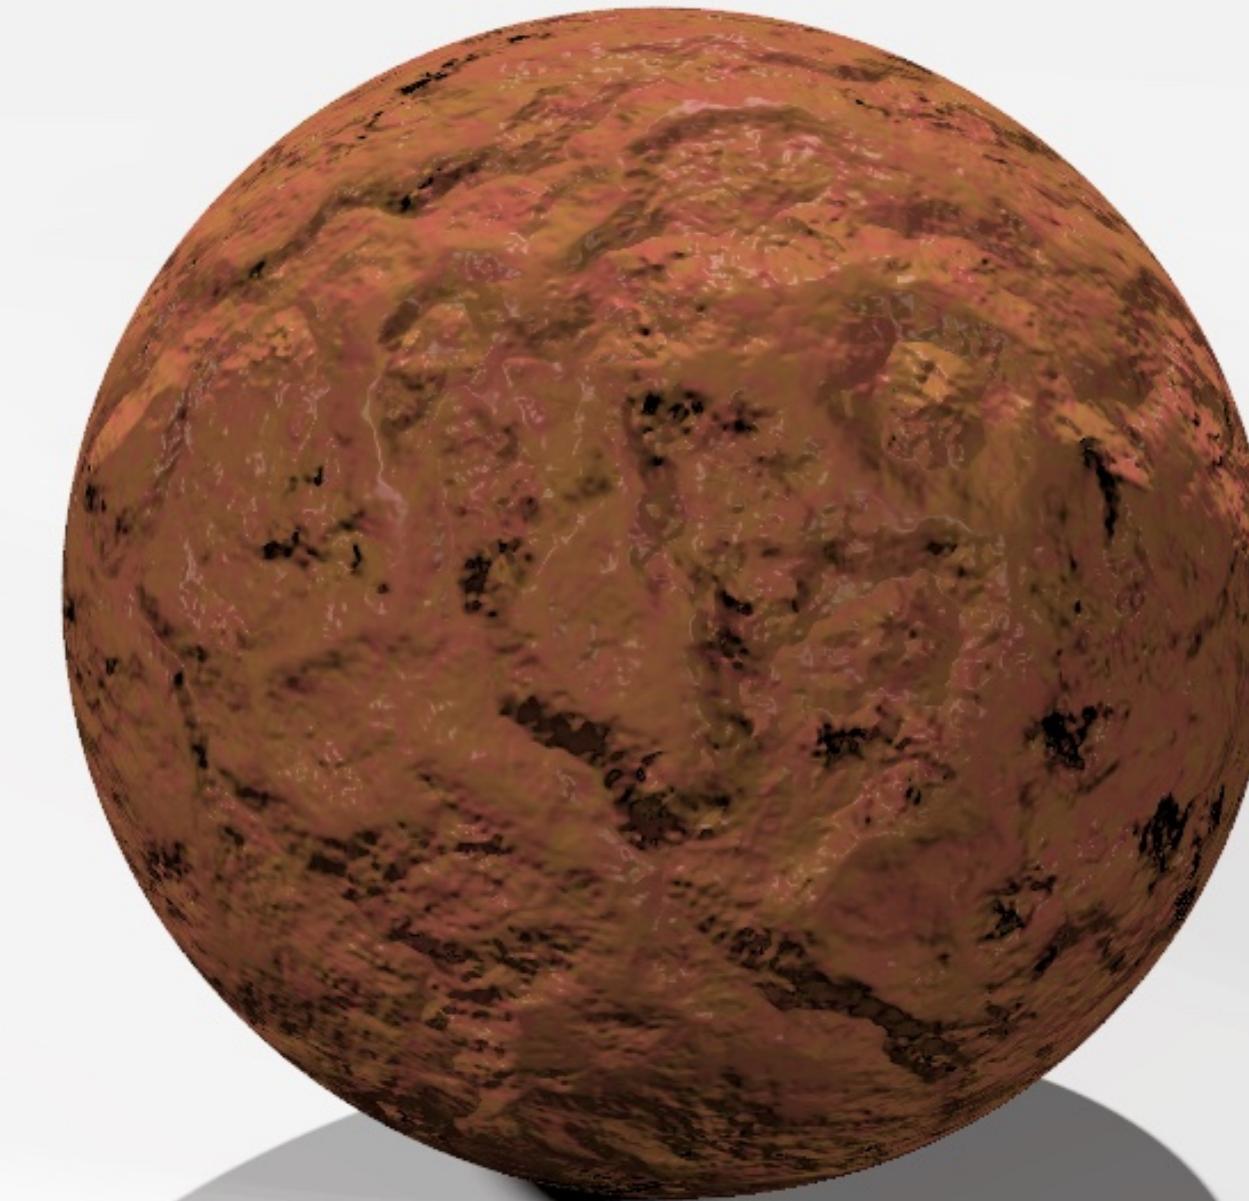
Stanford CS248A, Winter 2026

Layered material



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 vertex positions according to texture values
(note bumpy silhouette and shadow boundary)

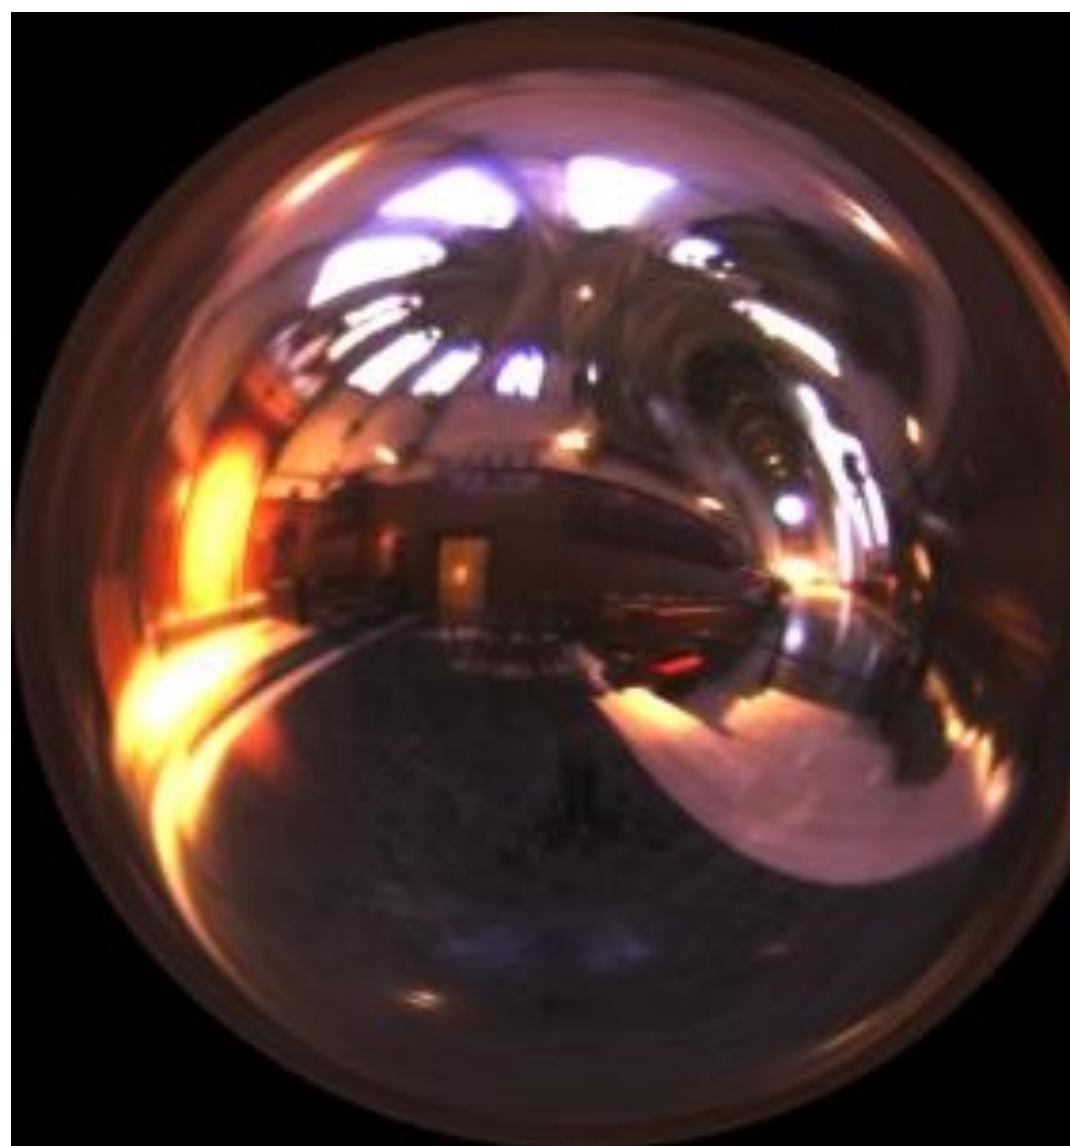
Represent precomputed lighting and shadows



Original model

With ambient occlusion

Extracted ambient occlusion map



Grace Cathedral environment map



Environment map used in a rendering

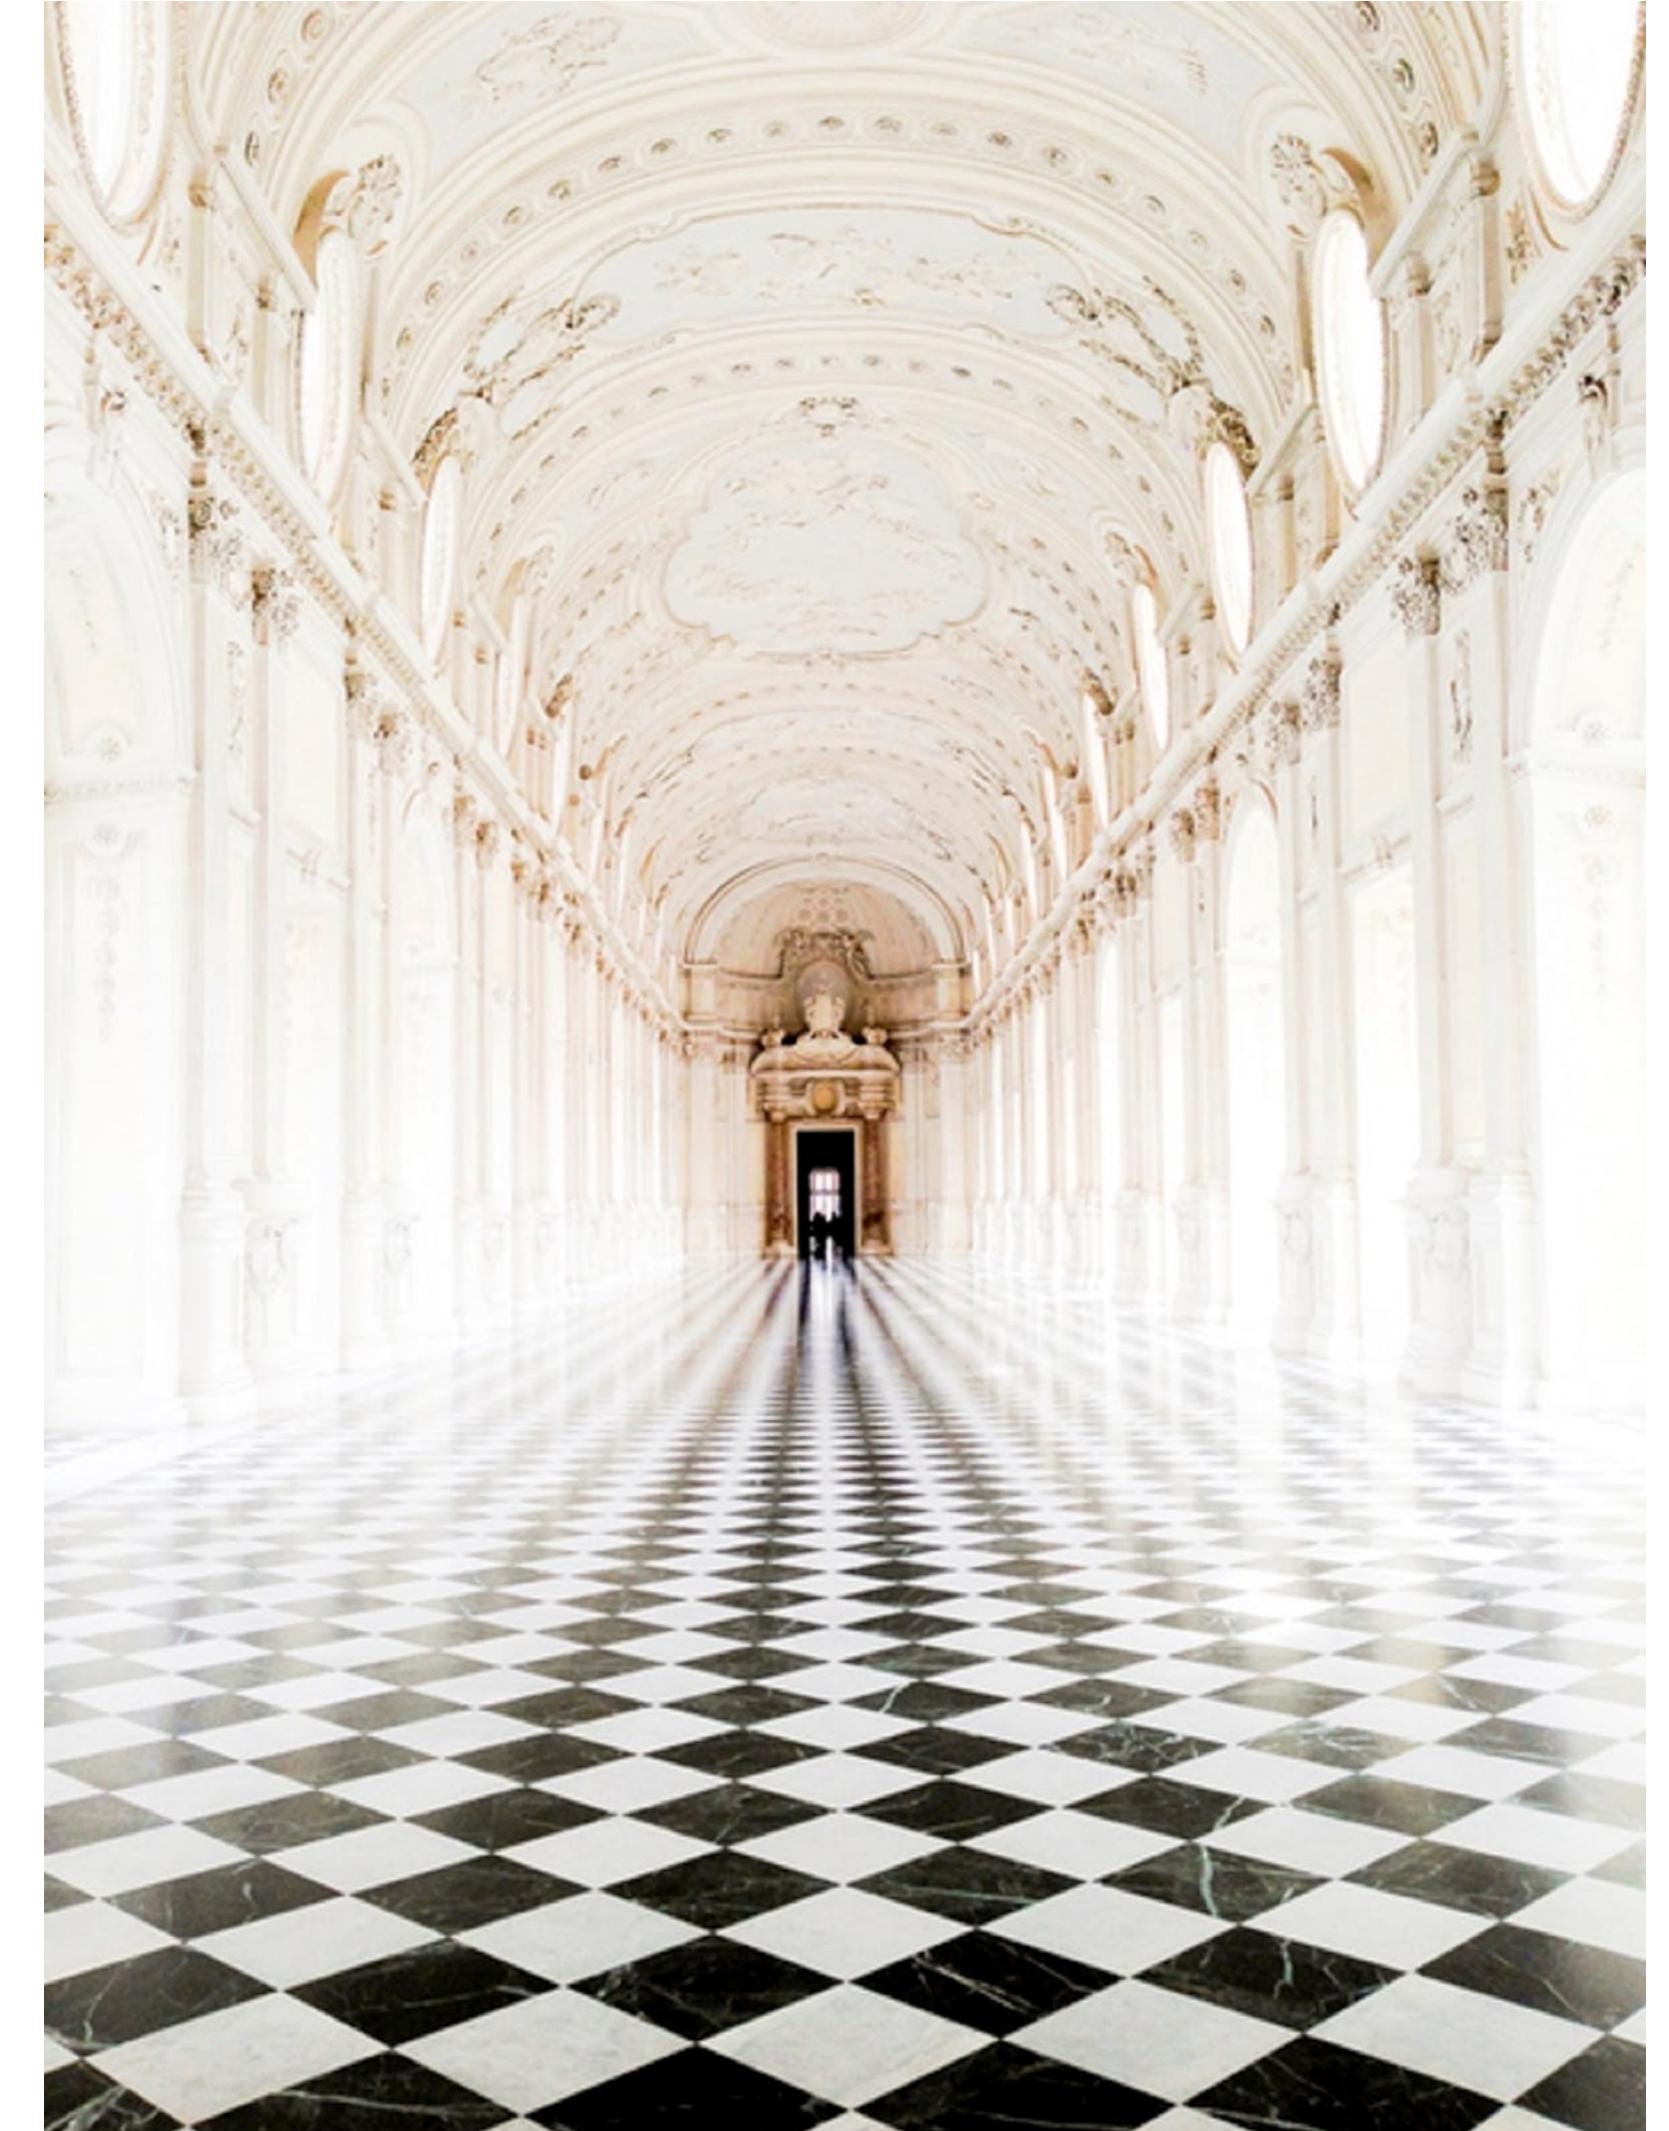
Perspective and texture

■ PREVIOUSLY:

- *transformations* (how to manipulate geometry in space)
- *rasterization/raytracings* (how to sample geometry into colored pixels)

■ TODAY:

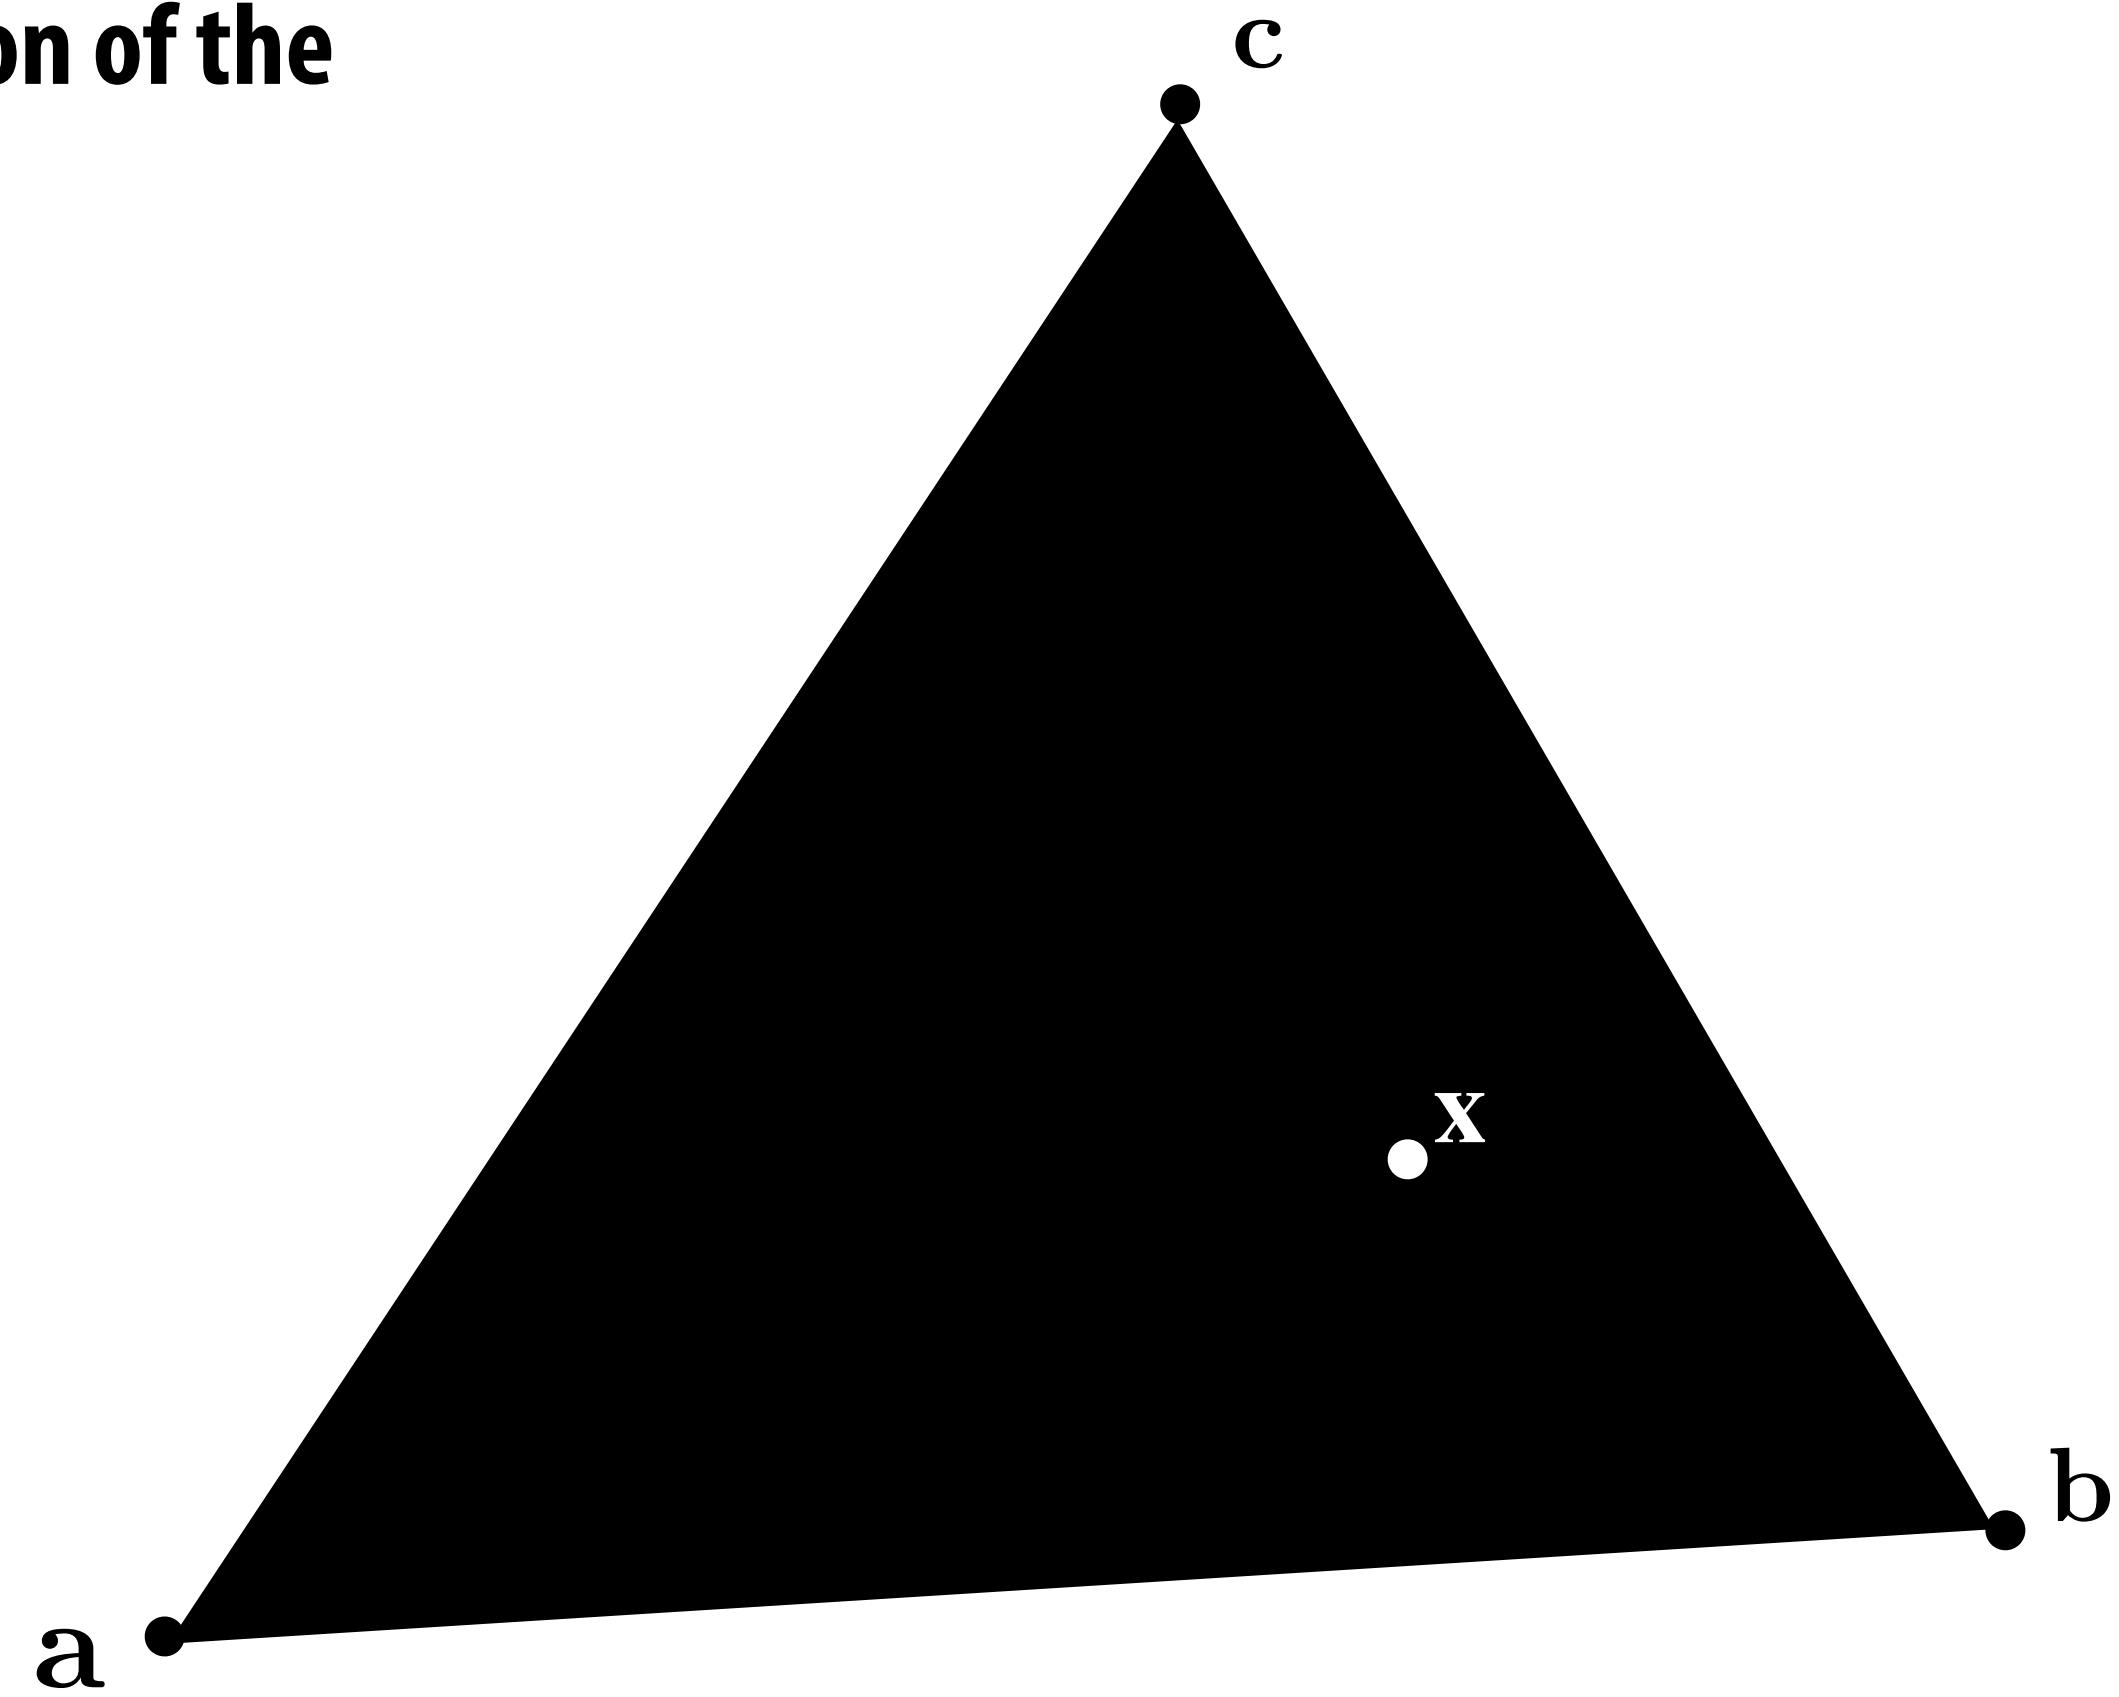
- **see where these two ideas come crashing together!**
- **talk about how to map *texture* onto a surface to get more detail**



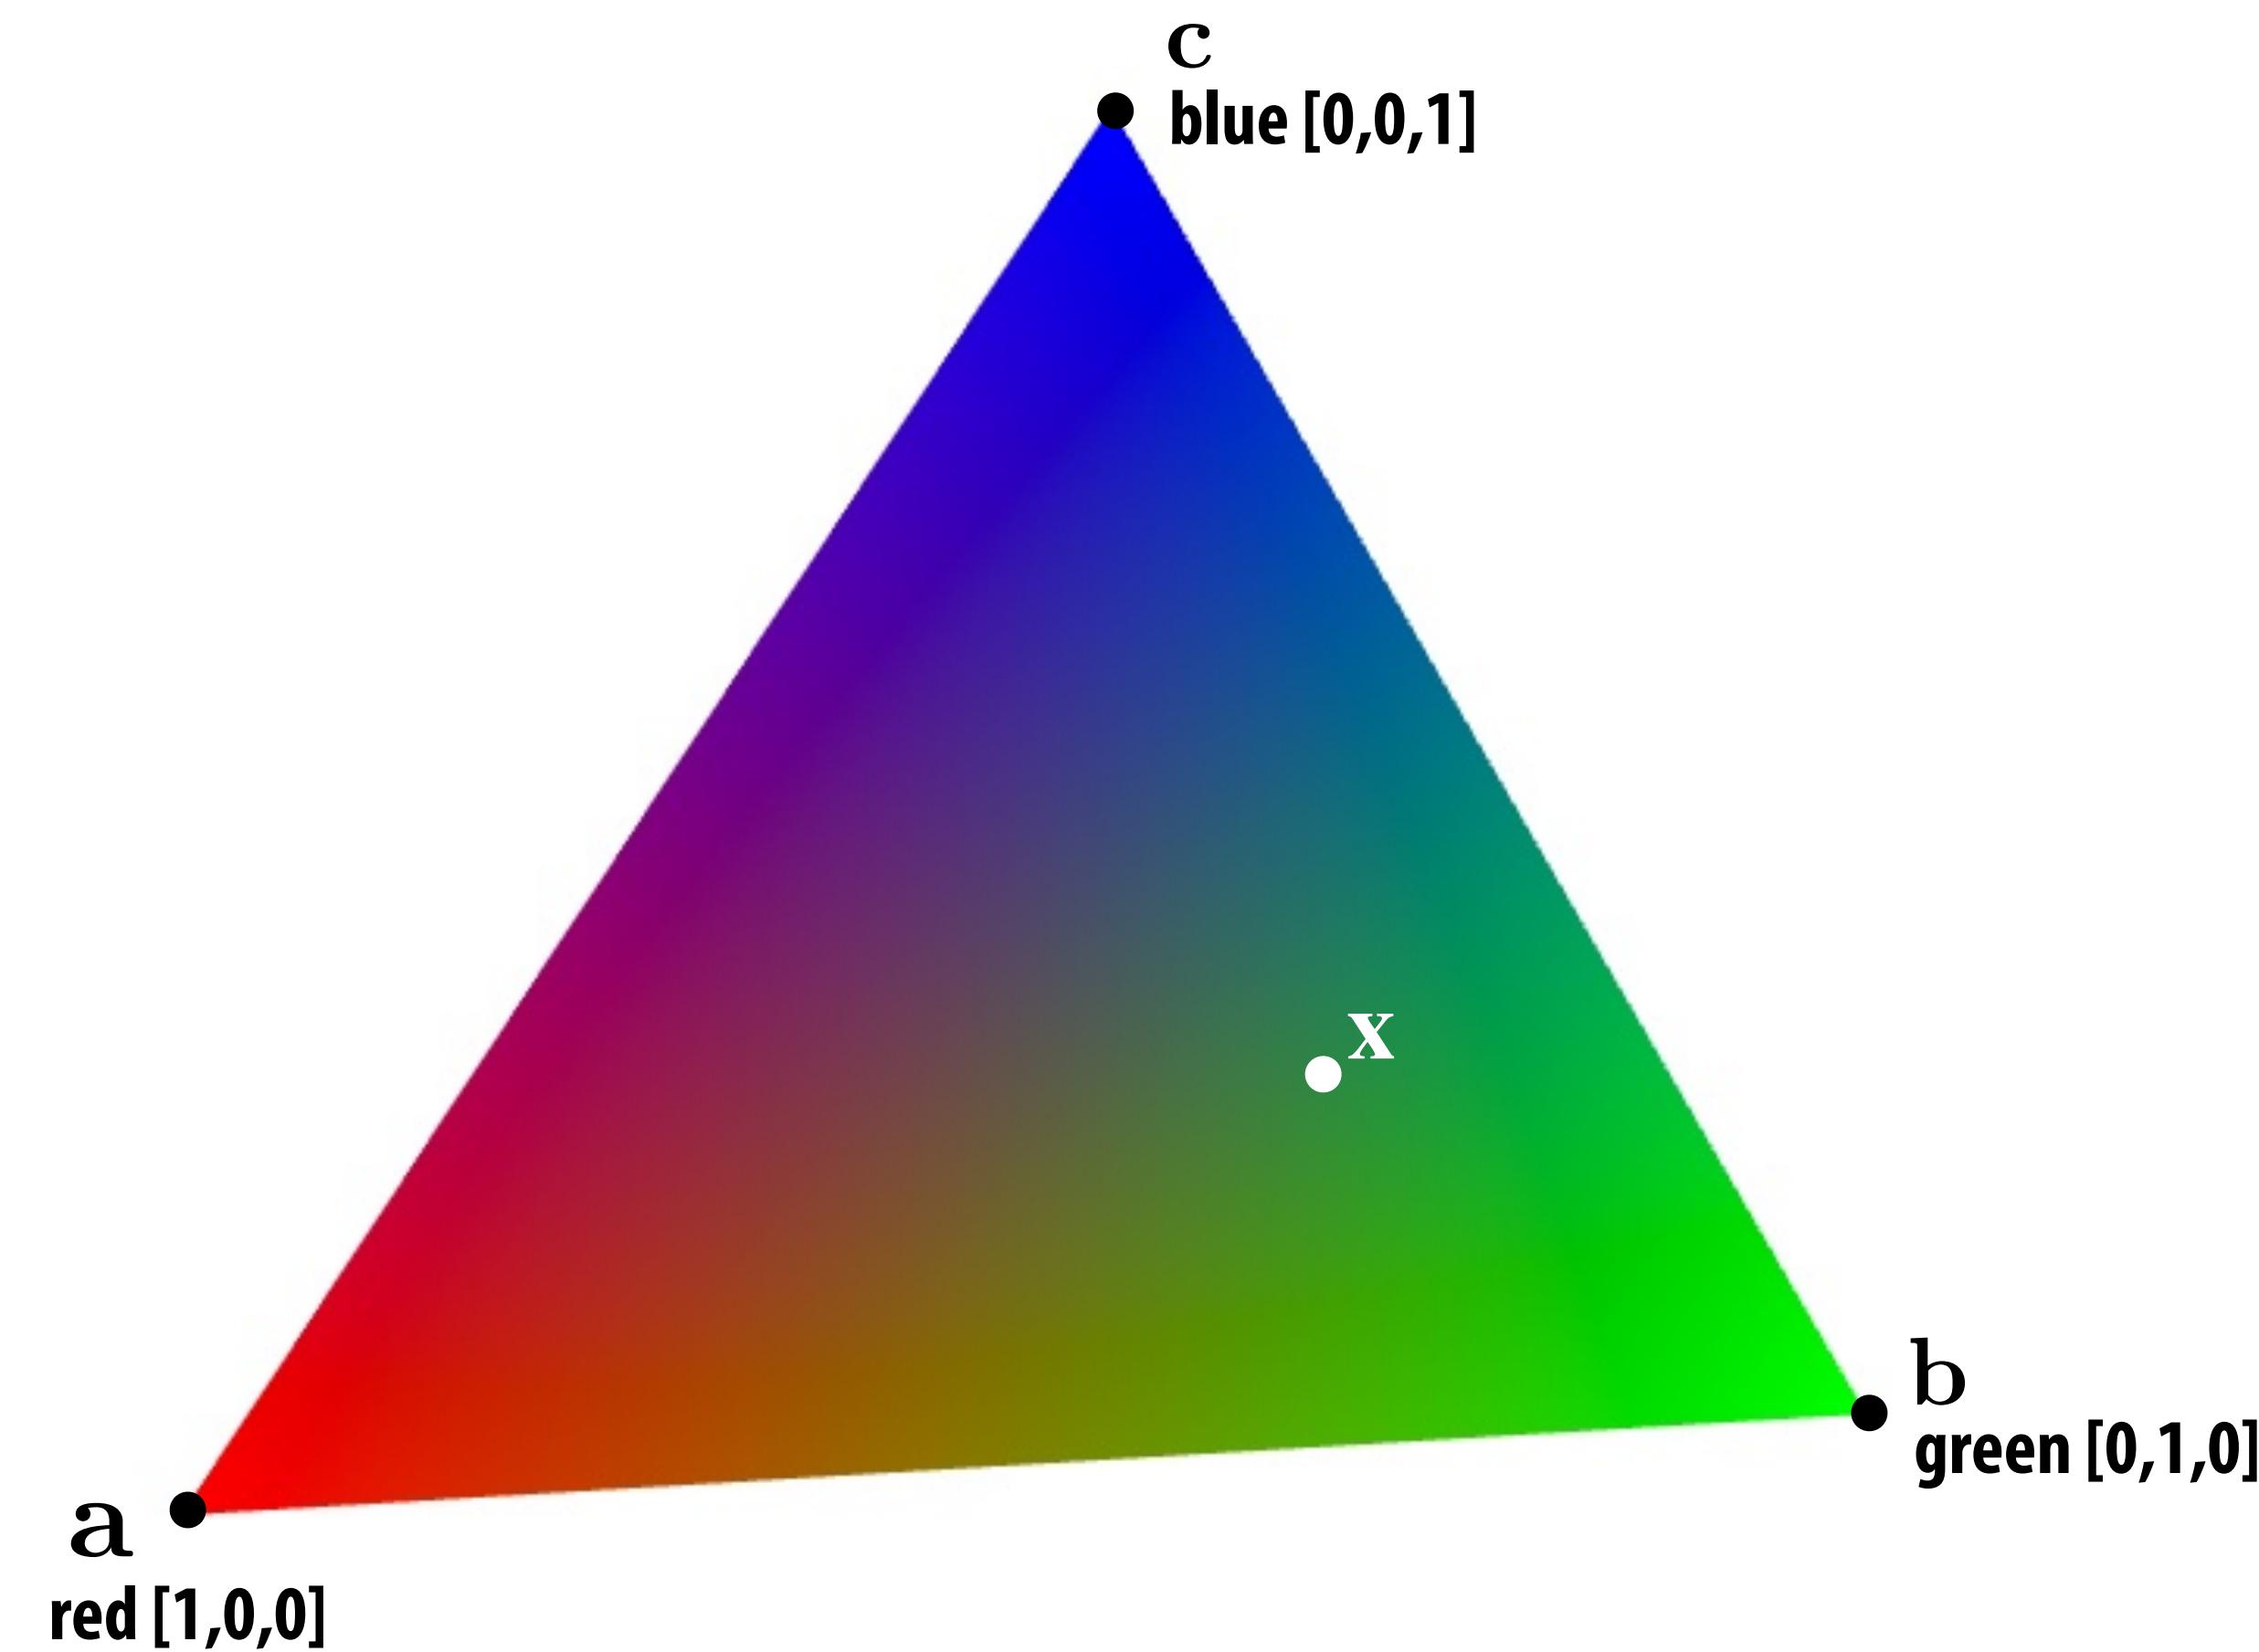
**Why is it hard to render
an image like this?**

Recall the function $\text{coverage}(x, y)$ from lecture 2

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



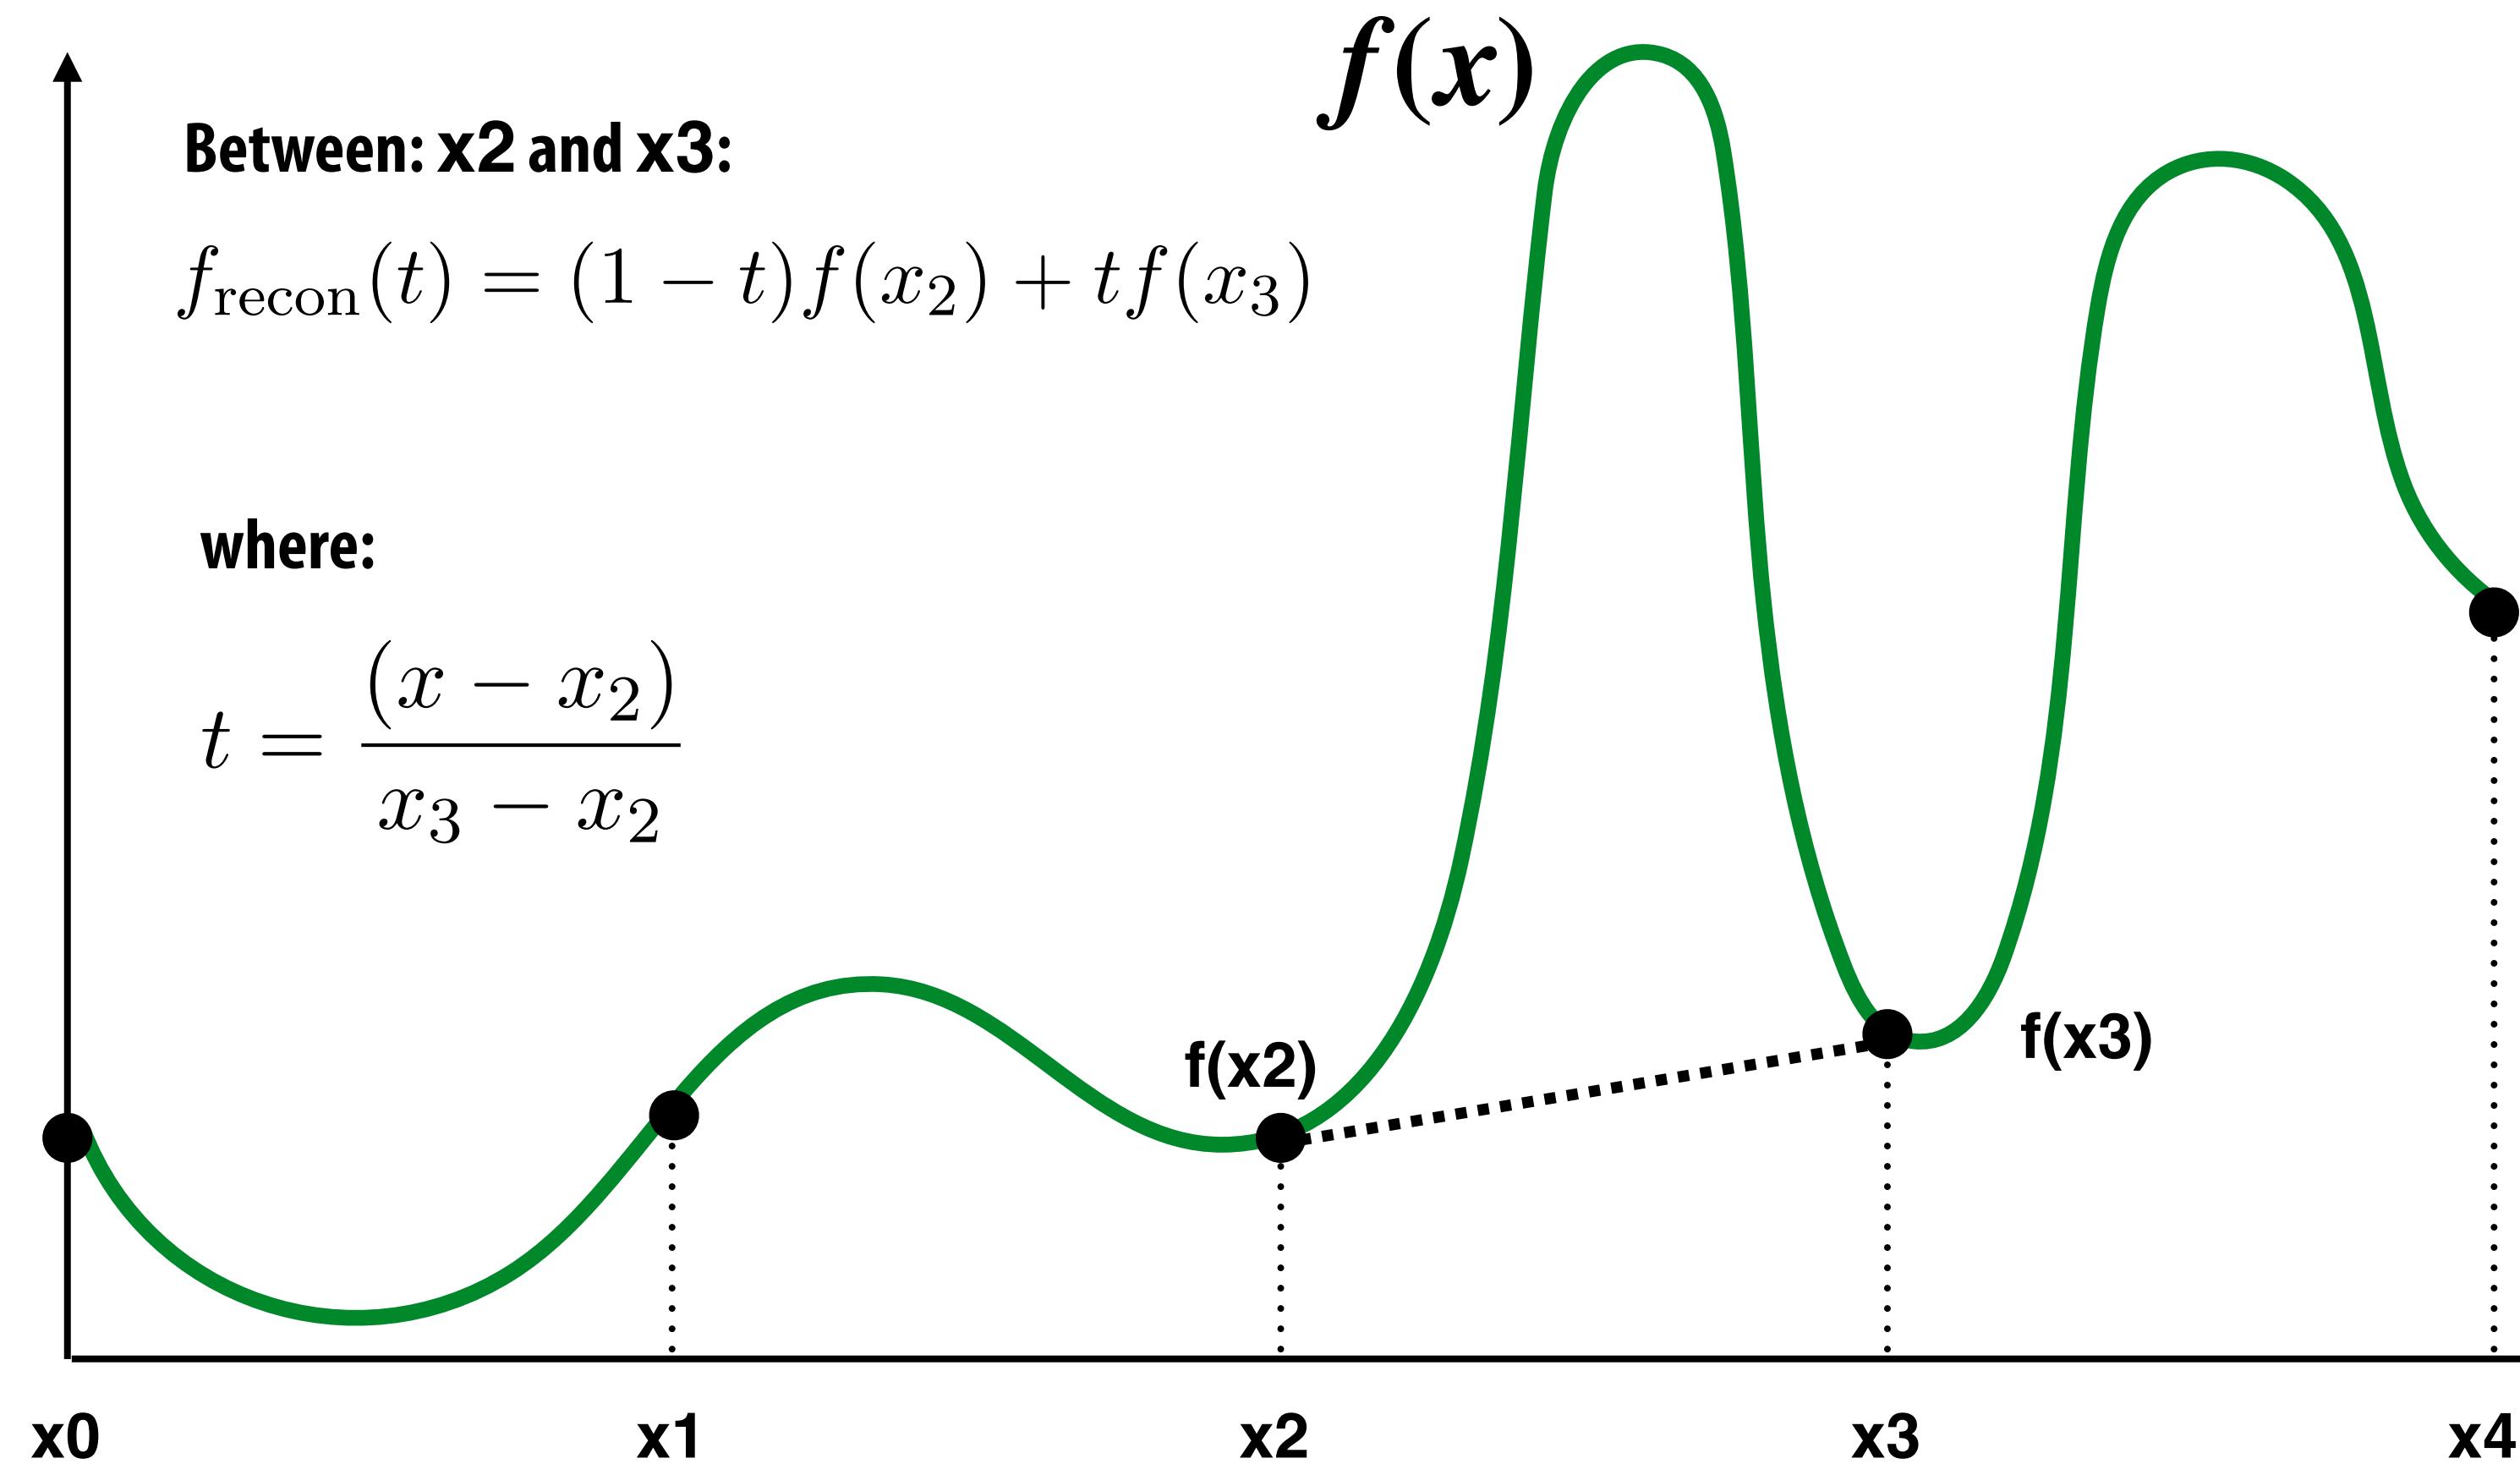
Consider sampling a different signal: $\text{color}(x,y)$



What is the triangle's color at the point **x** , given its colors at points **a**, **b**, **c** ?

Review: interpolation in 1D

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

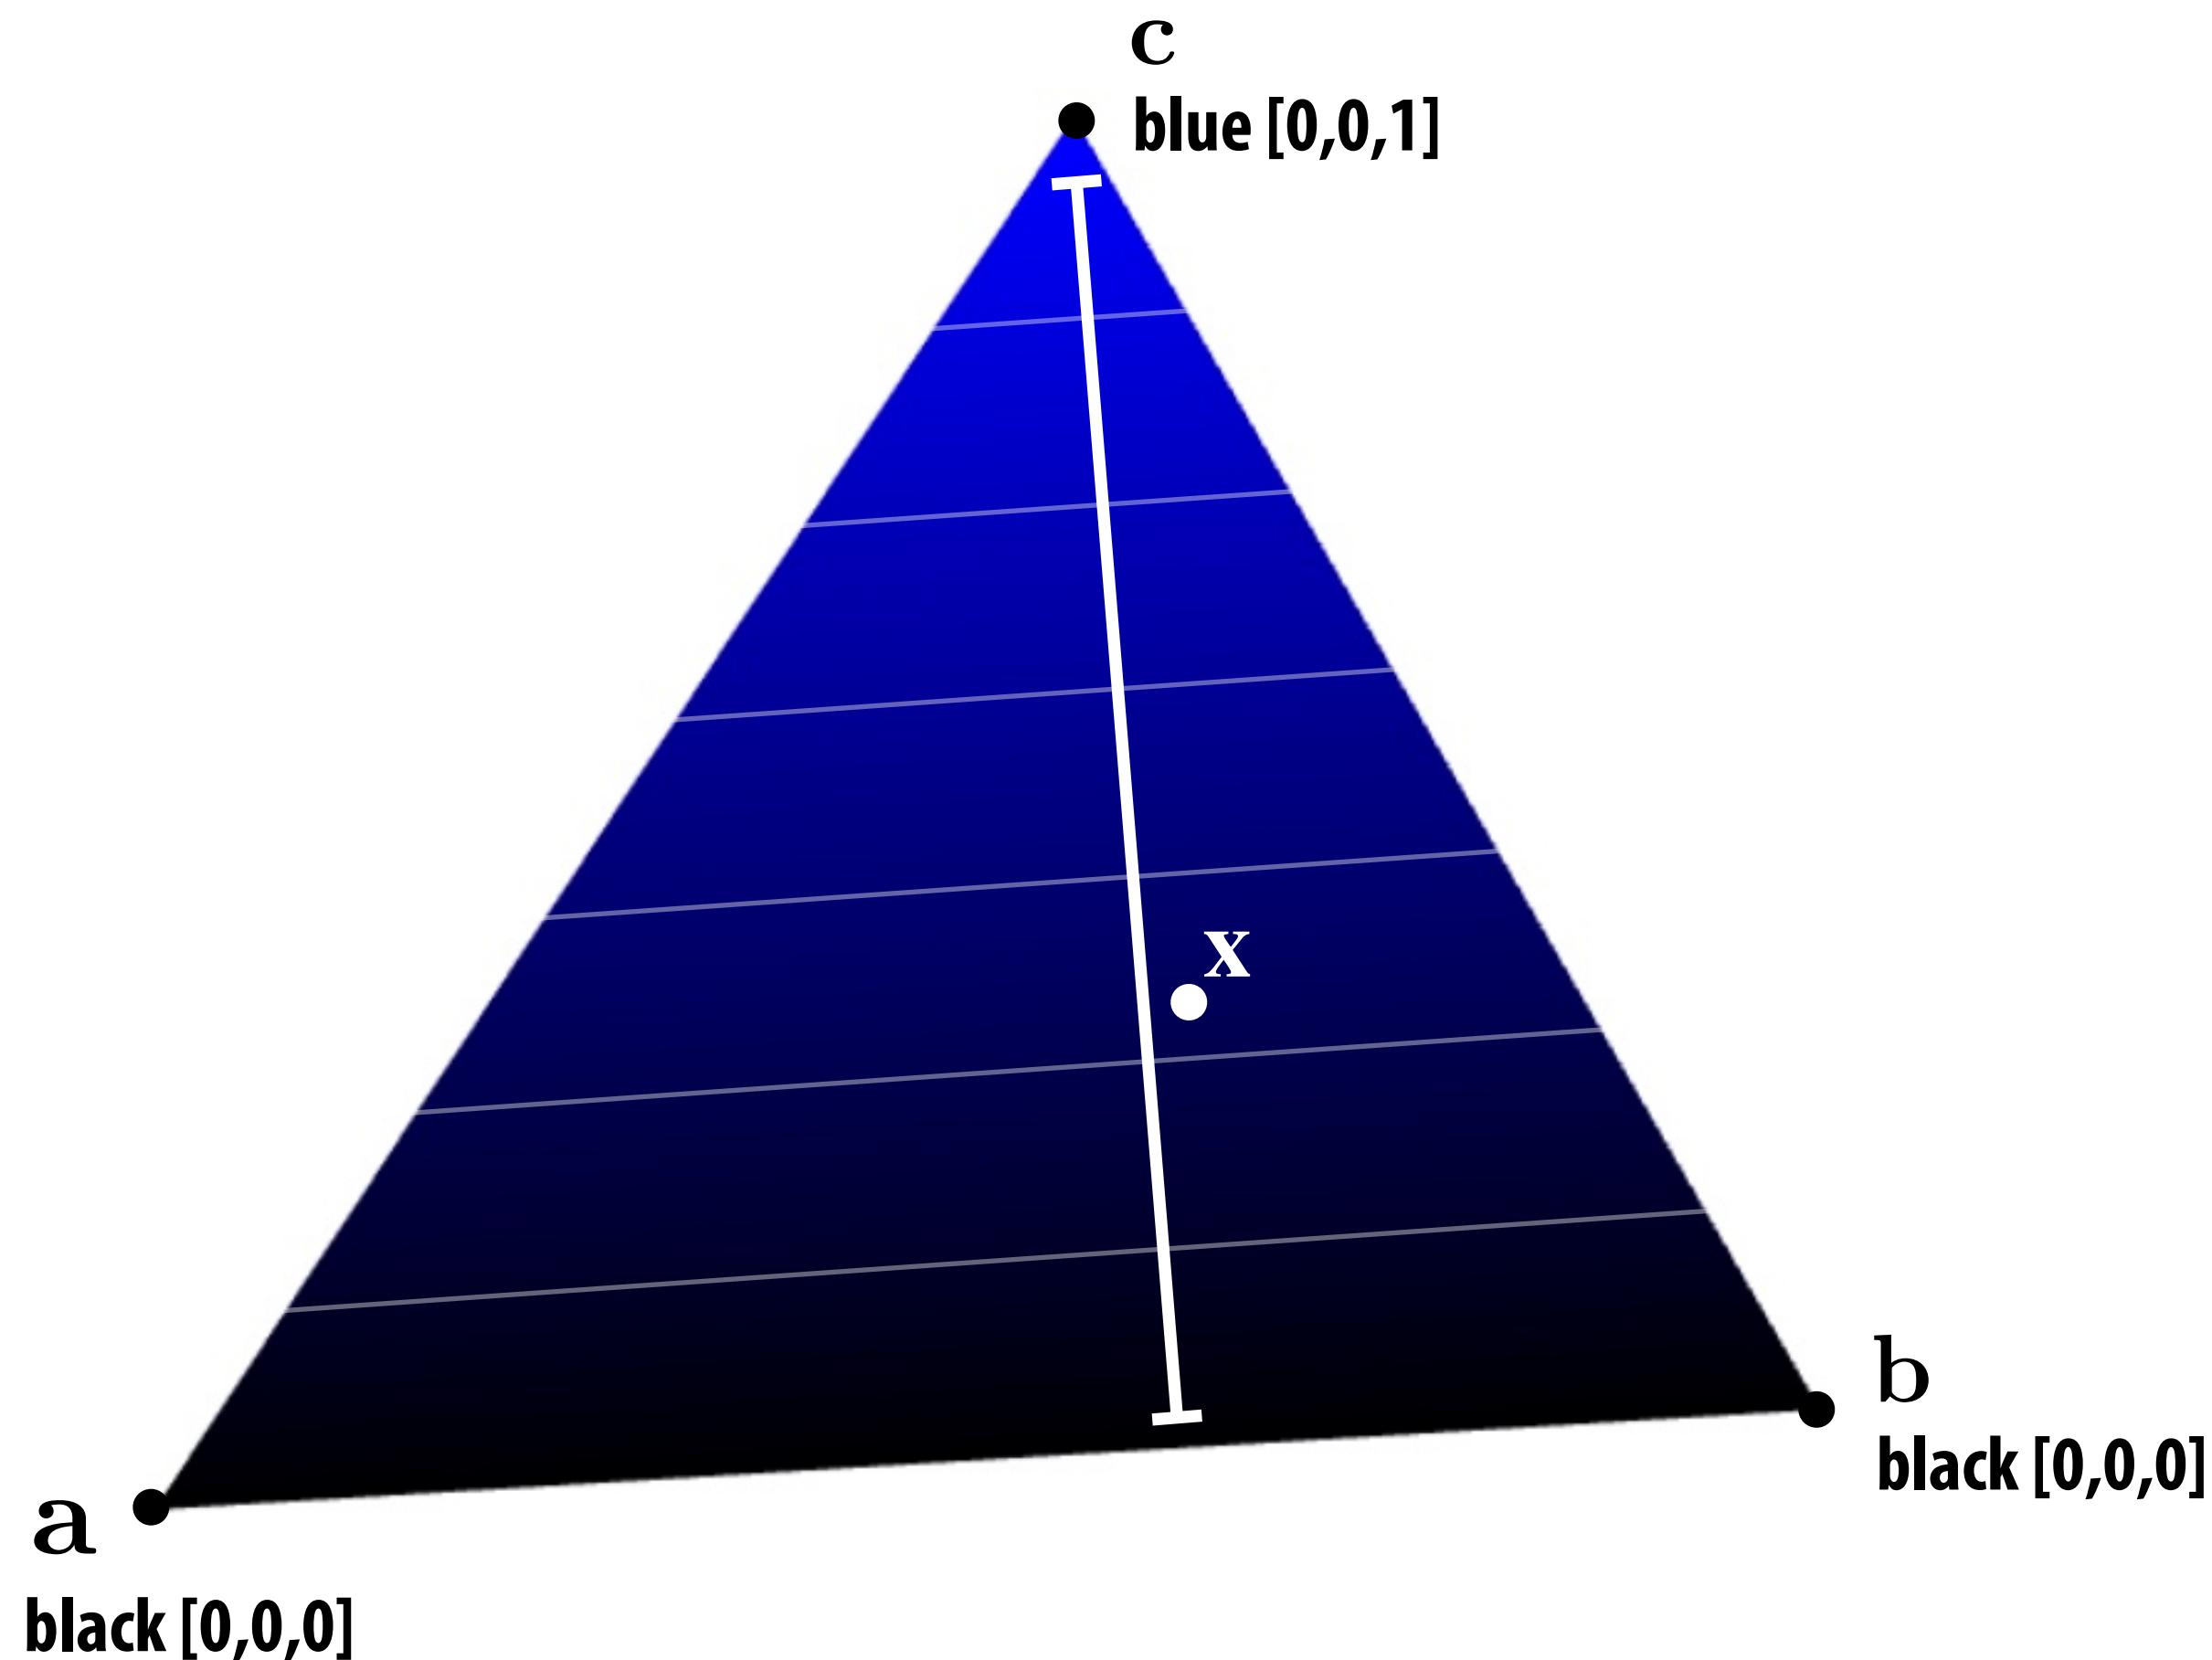


Review: consider similar behavior on triangle

Color depends on distance from $b - a$

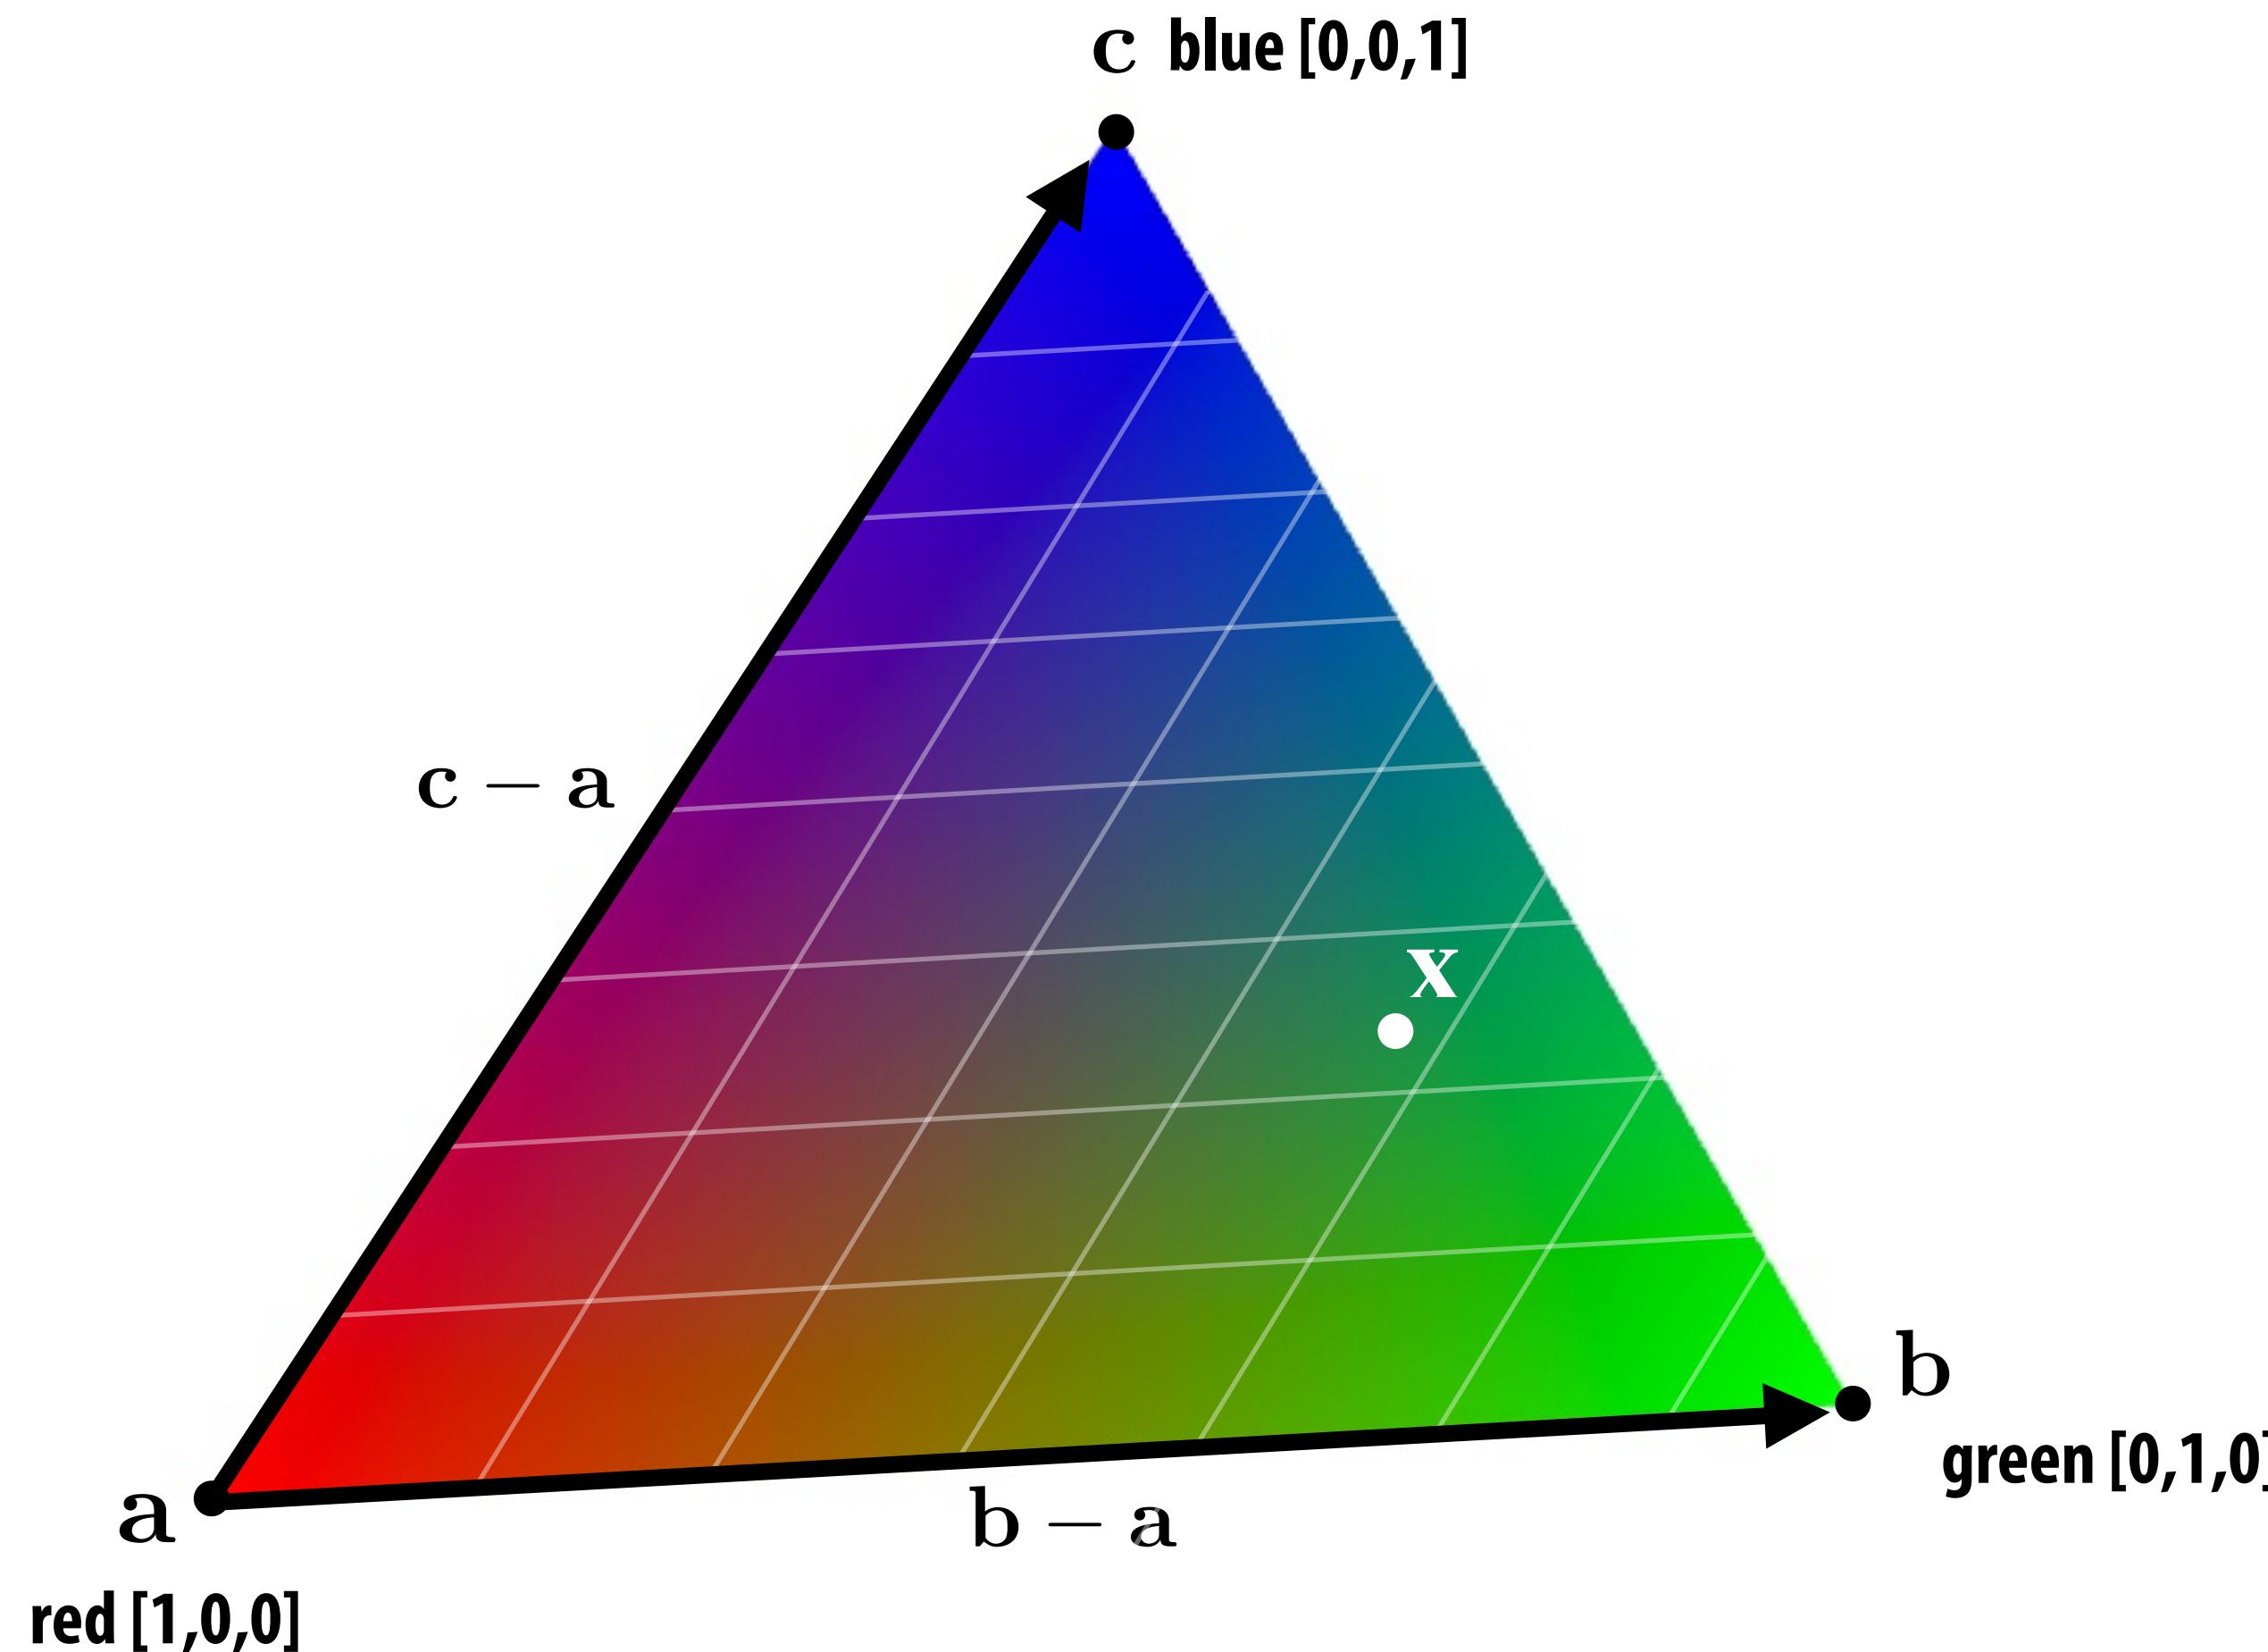
$$\text{Color} = (1 - t) [0 \ 0 \ 0] + t [0 \ 0 \ 1]$$

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



How can we interpolate in 2D between three values?

Review: linear interpolation of quantities over triangle



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

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

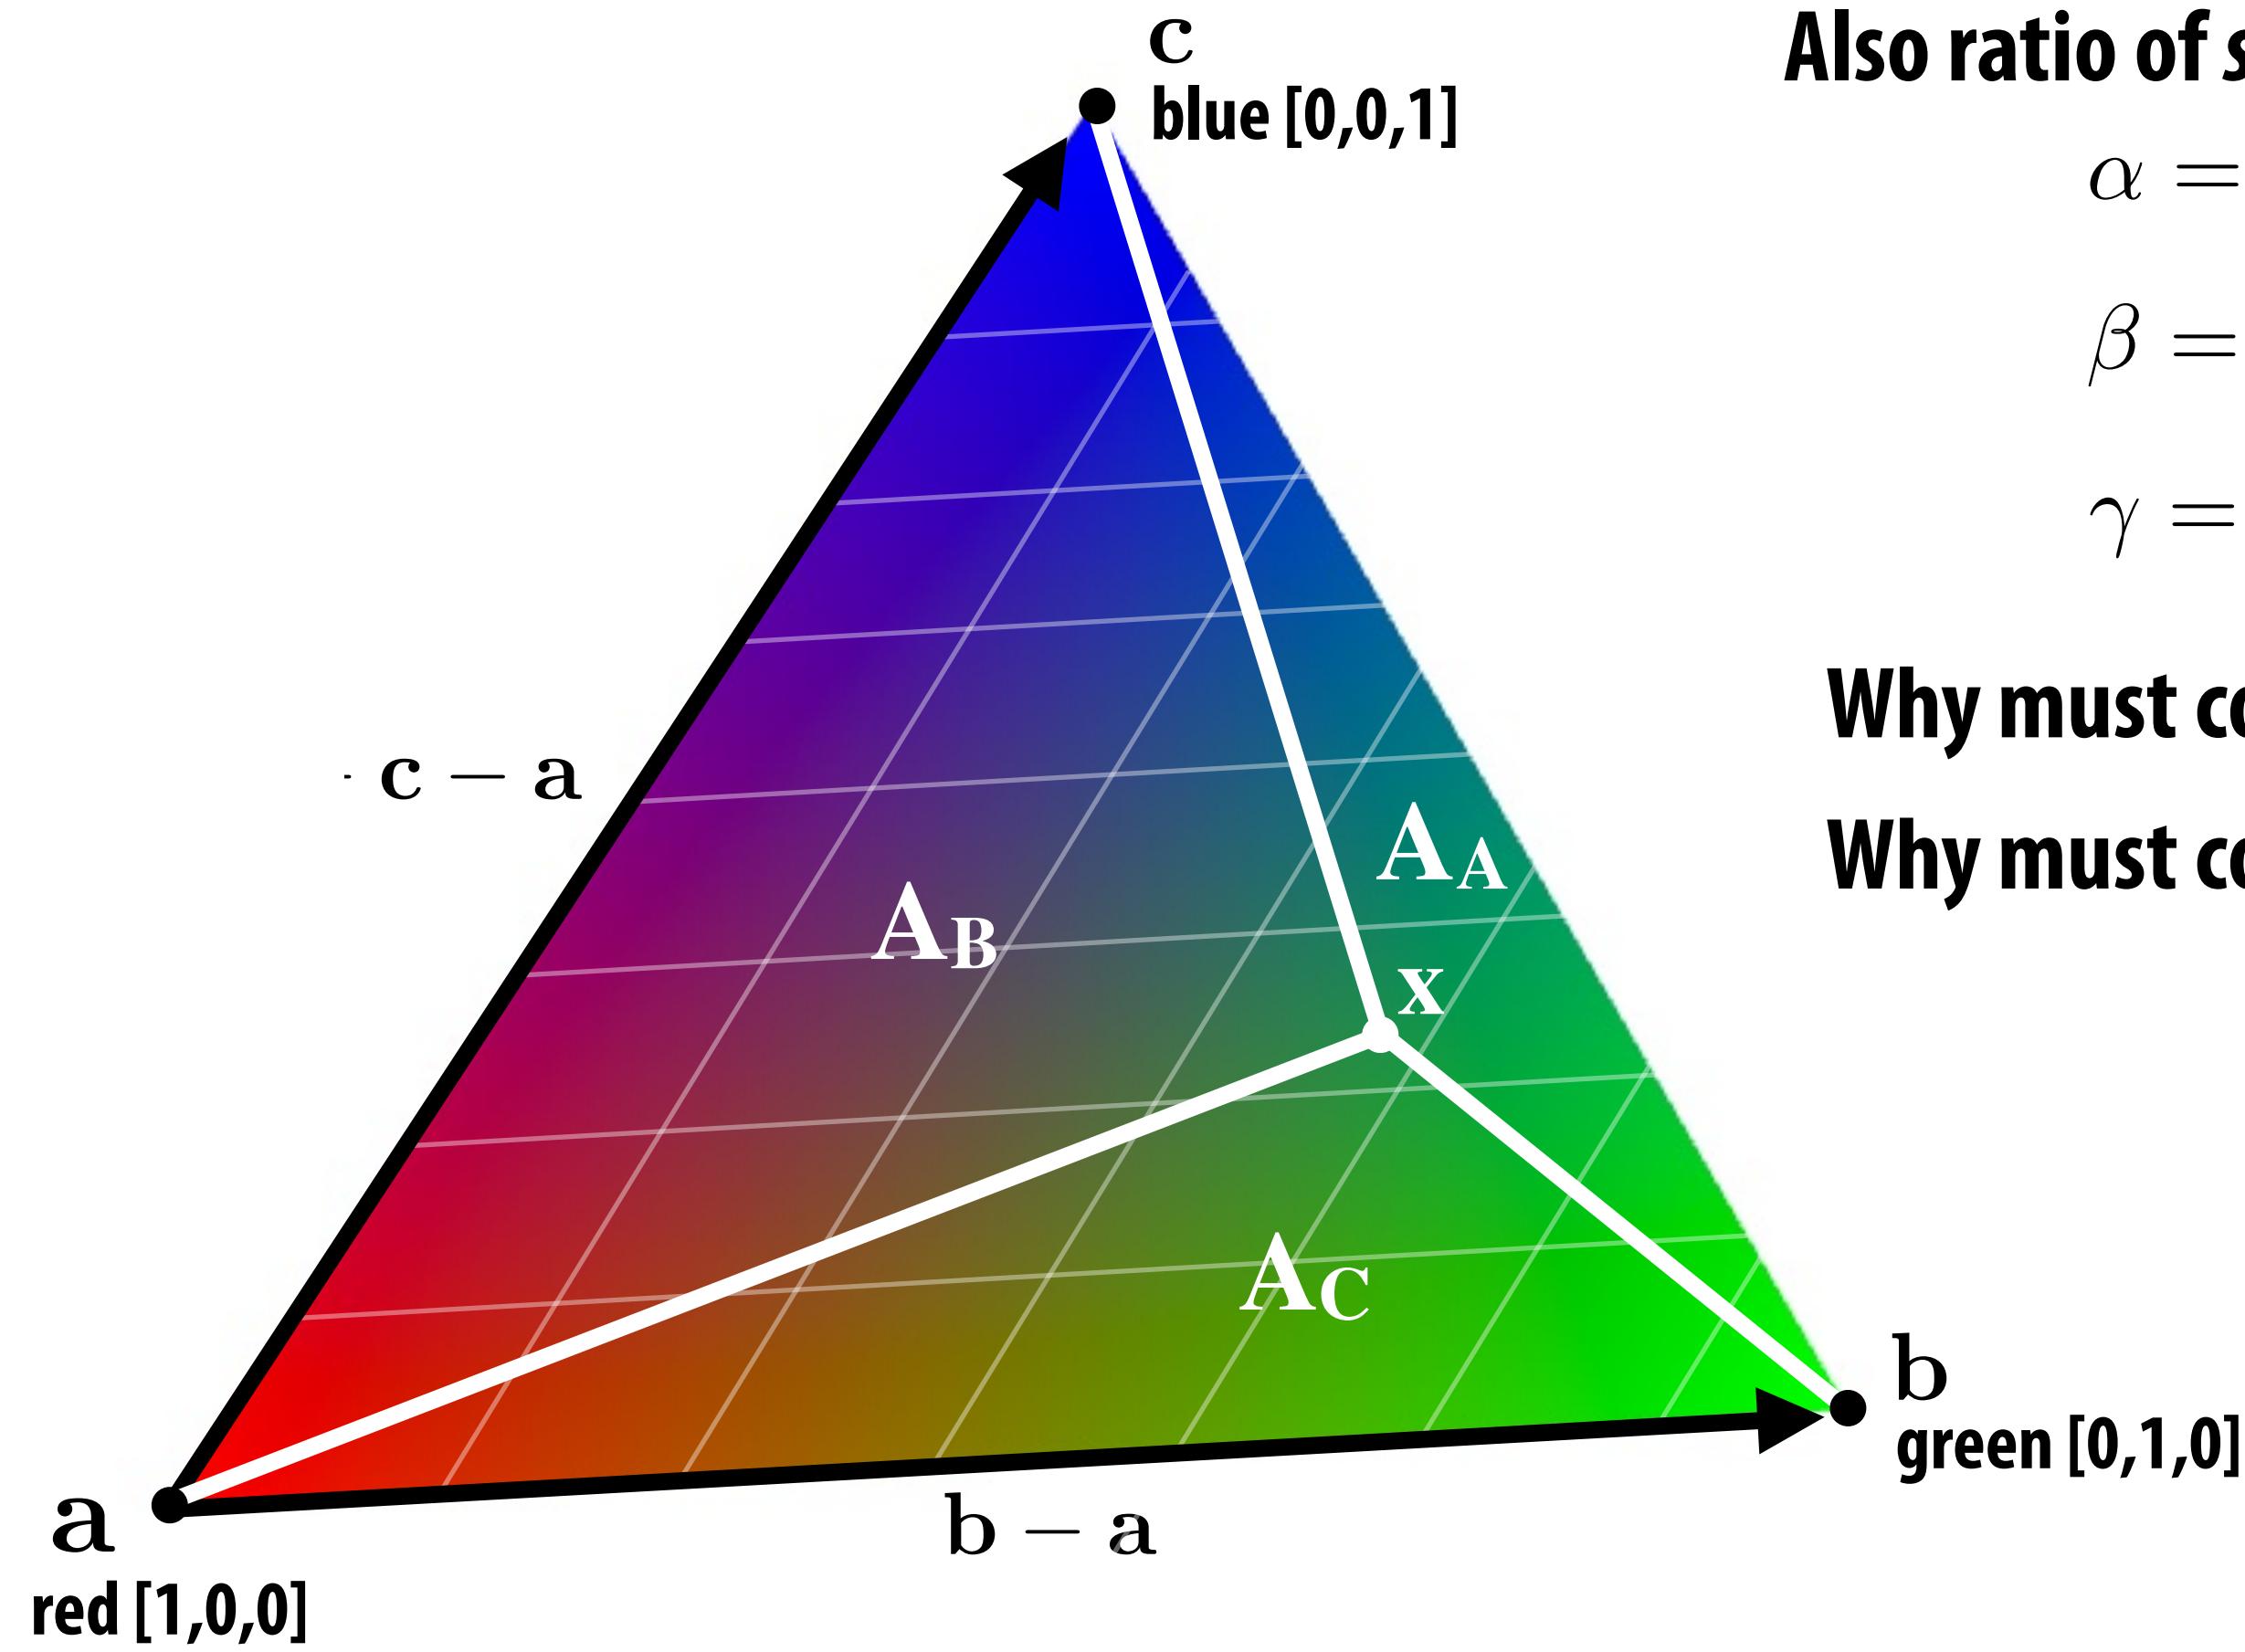
$$\alpha + \beta + \gamma = 1$$

Color value at \mathbf{X} is linear combination of color value at three triangle vertices.

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

What is the triangle's color at the point \mathbf{x} , given its color at points \mathbf{a} , \mathbf{b} , \mathbf{c} ?

Another way: barycentric coordinates as ratio of areas



Also ratio of *signed* areas:

$$\alpha = A_A/A$$

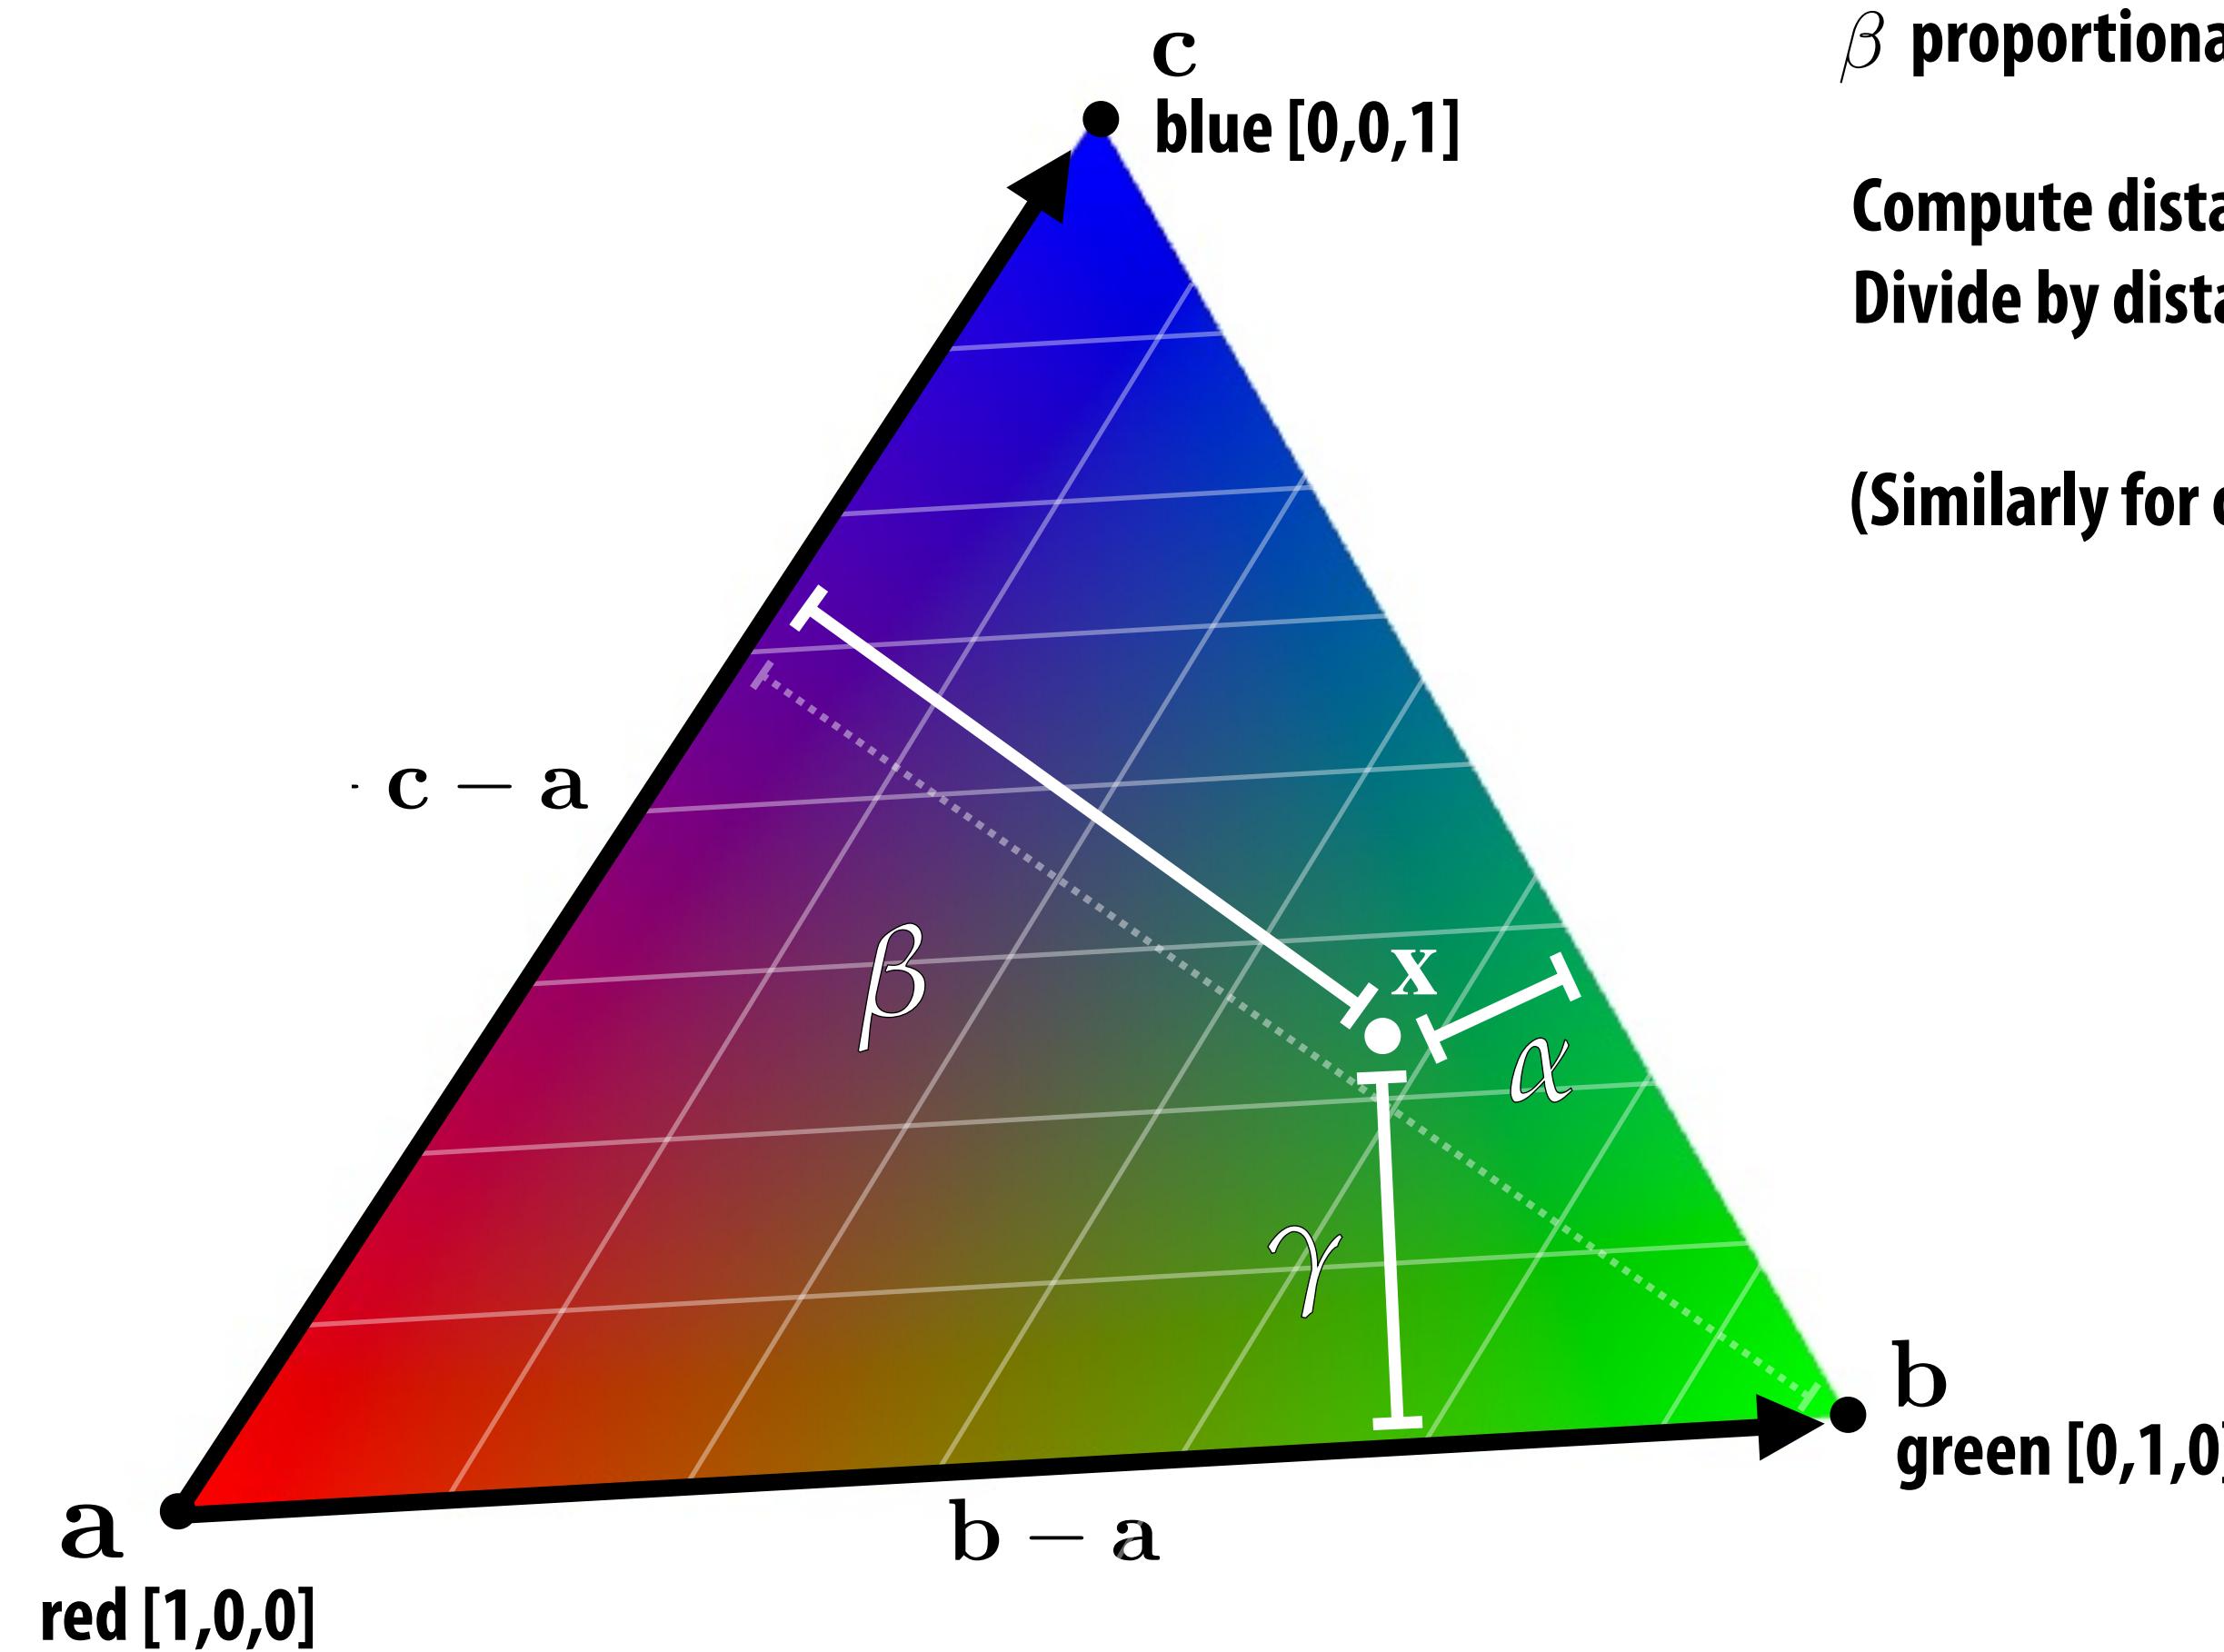
$$\beta = A_B/A$$

$$\gamma = A_C/A$$

Why must coordinates sum to one?

Why must coordinates be between 0 and 1?

Yet another way: barycentric coordinates as scaled distances



β proportional to distance from **X** to edge **c — a**

Compute distance of **X** from line **c — a**

Divide by distance of **b** from line **c — a** ("height" of triangle)

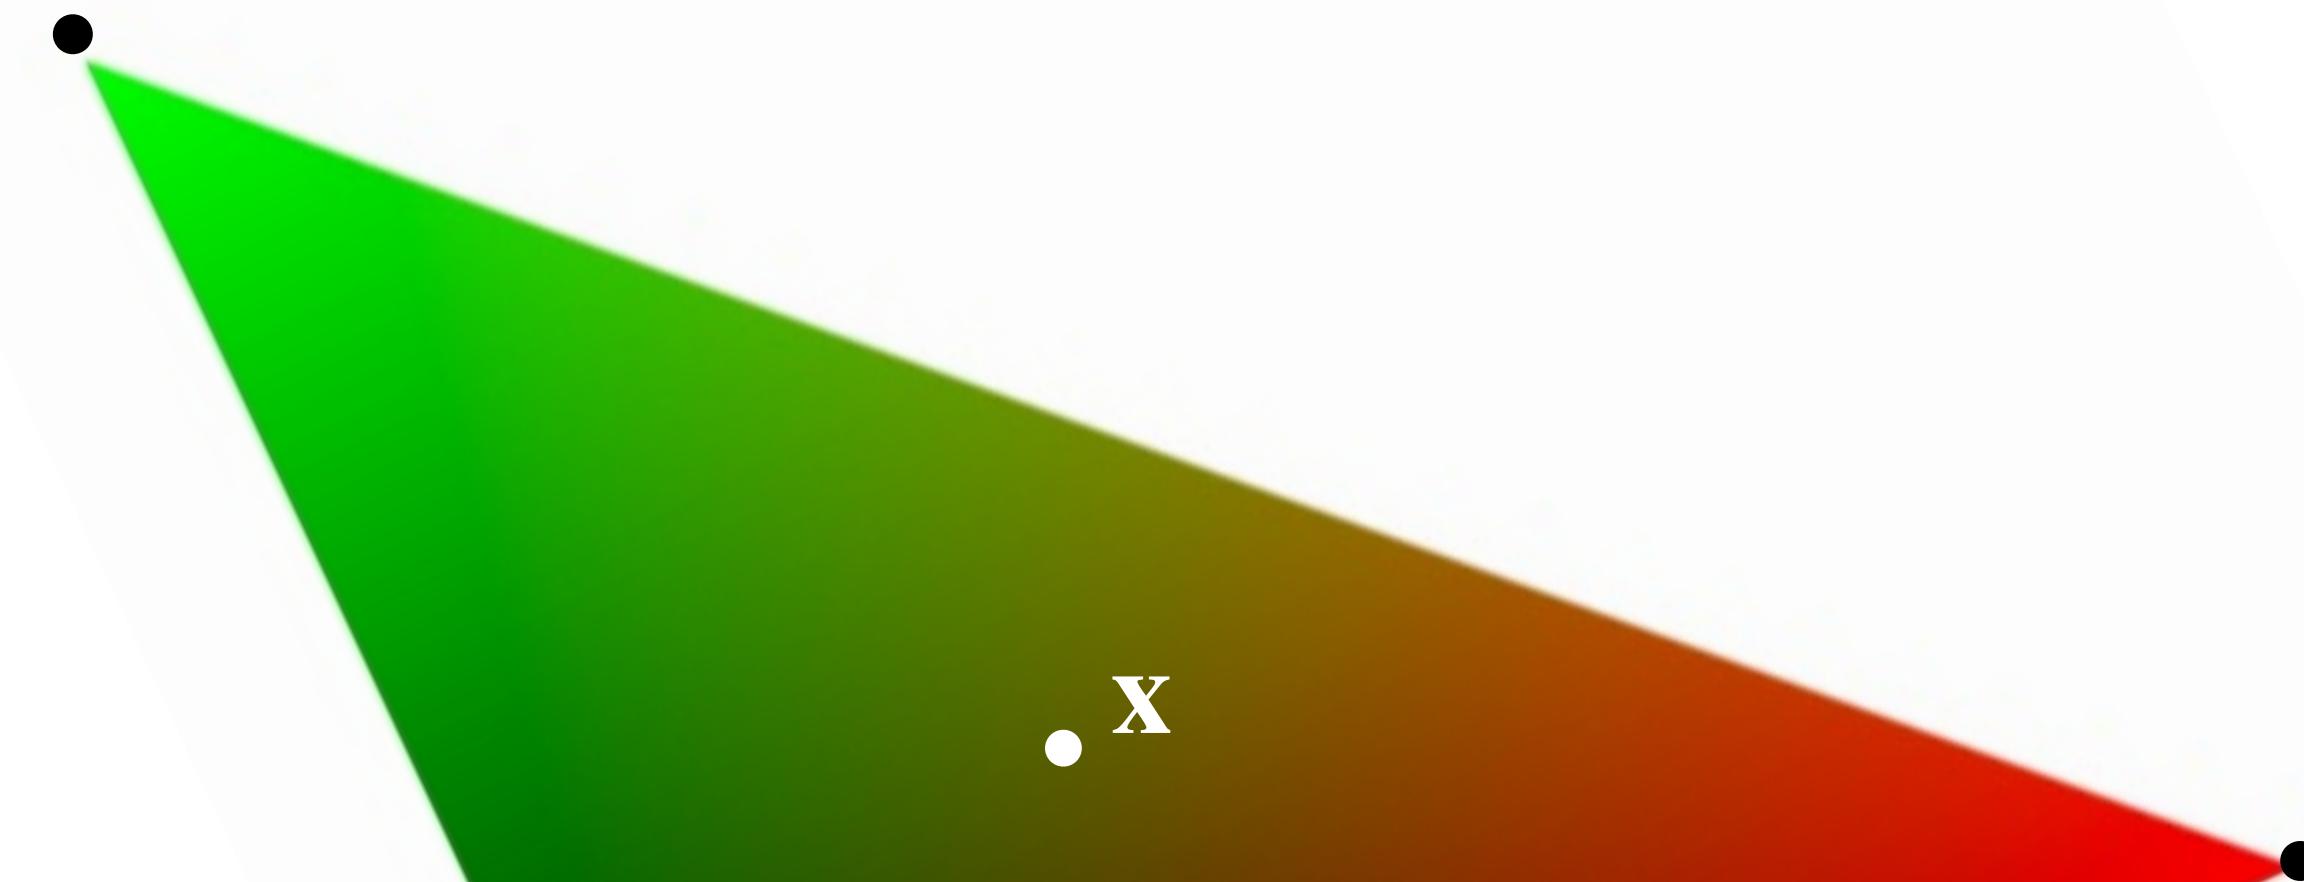
(Similarly for other two barycentric coordinates)

You can linearly interpolate any values (defined at vertices) over the triangle this way

Here, I'm interpolating position (x, y, z), color (r, g, b), and extra values (u, v)

$$C = (x_2, y_2, z_2, r_2, g_2, b_2, u_2, v_2)$$

Vertex is green, so $(r_2, g_2, b_2) = (0, 1, 0)$



$$A = (x_0, y_0, z_0, r_0, g_0, b_0, u_0, v_0)$$

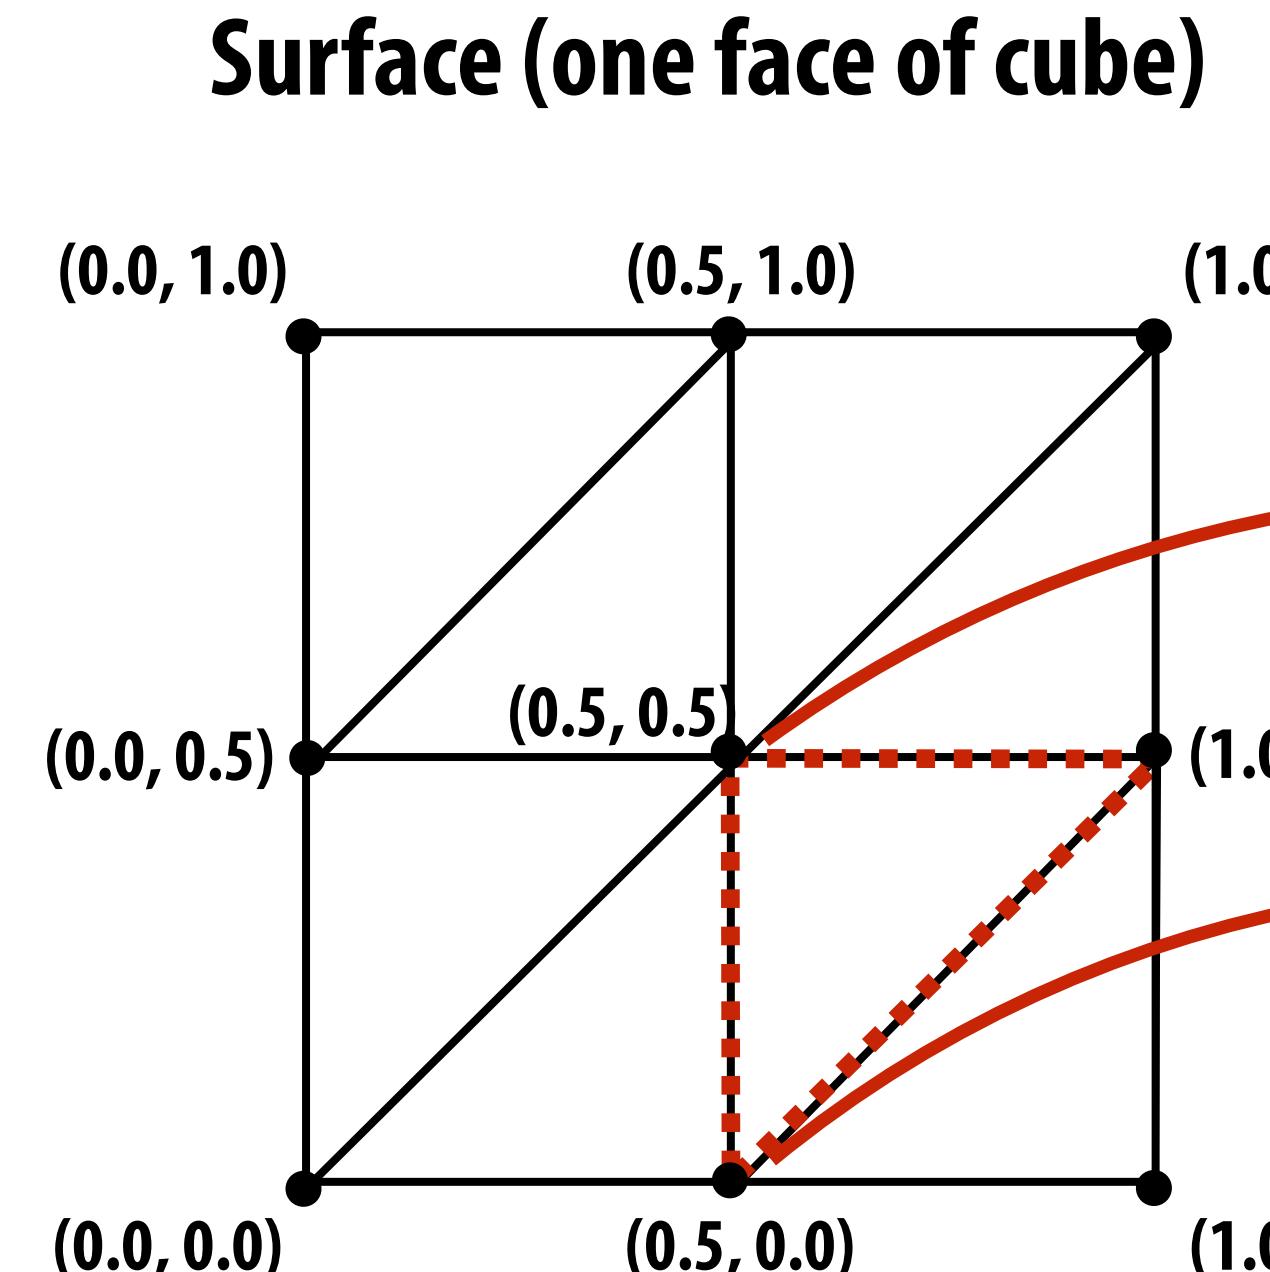
Vertex is black, so $(r_0, g_0, b_0) = (0, 0, 0)$

$$B = (x_1, y_1, z_1, r_1, g_1, b_1, u_1, v_1)$$

Vertex is red, so $(r_1, g_1, b_1) = (1, 0, 0)$

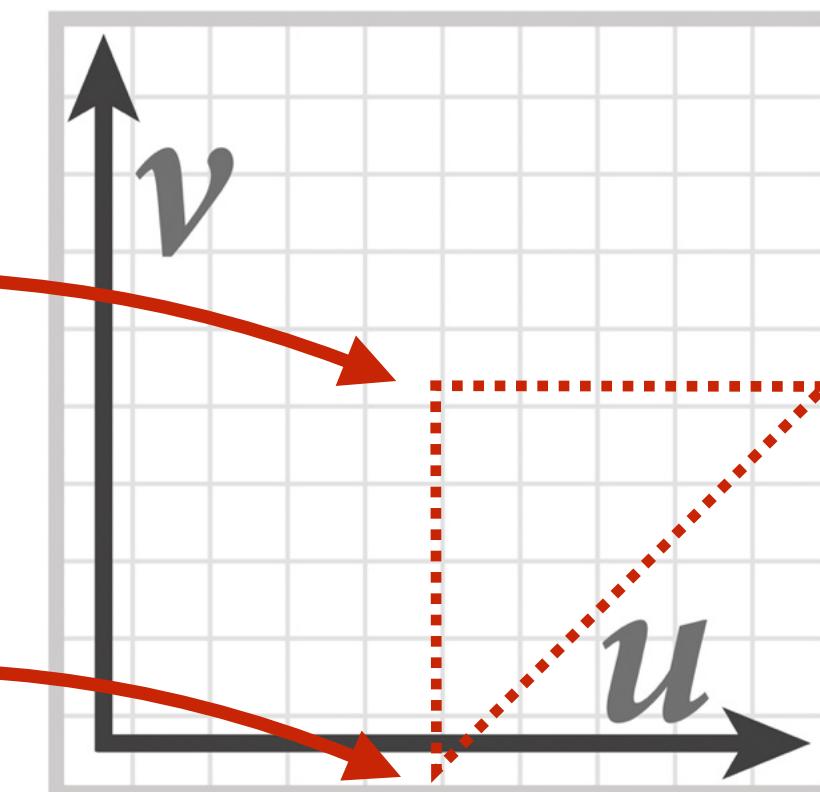
Texture coordinates

“Texture coordinates” define a mapping from surface coordinates (e.g., points on triangle) to points in the domain of a texture image



Eight triangles (one face of cube) with surface parameterization provided as per vertex texture coordinates (u, v)

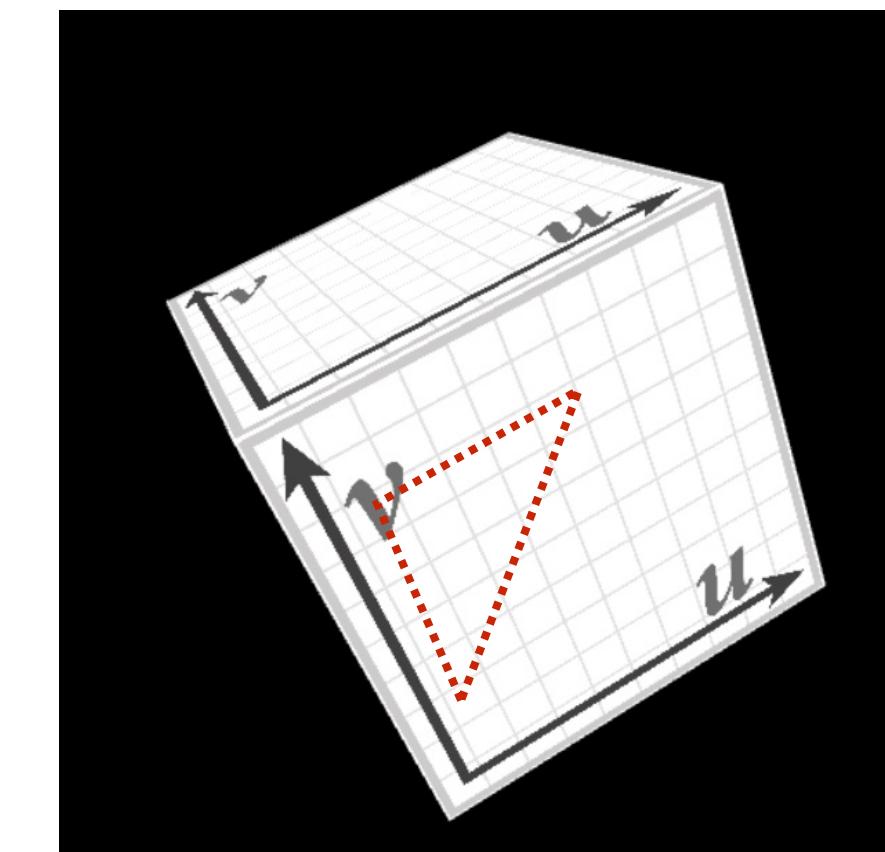
**Texture function
(represented by an image)**



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

Location of highlighted triangle in texture space shown in red.

Rendered image of texture mapped onto surface

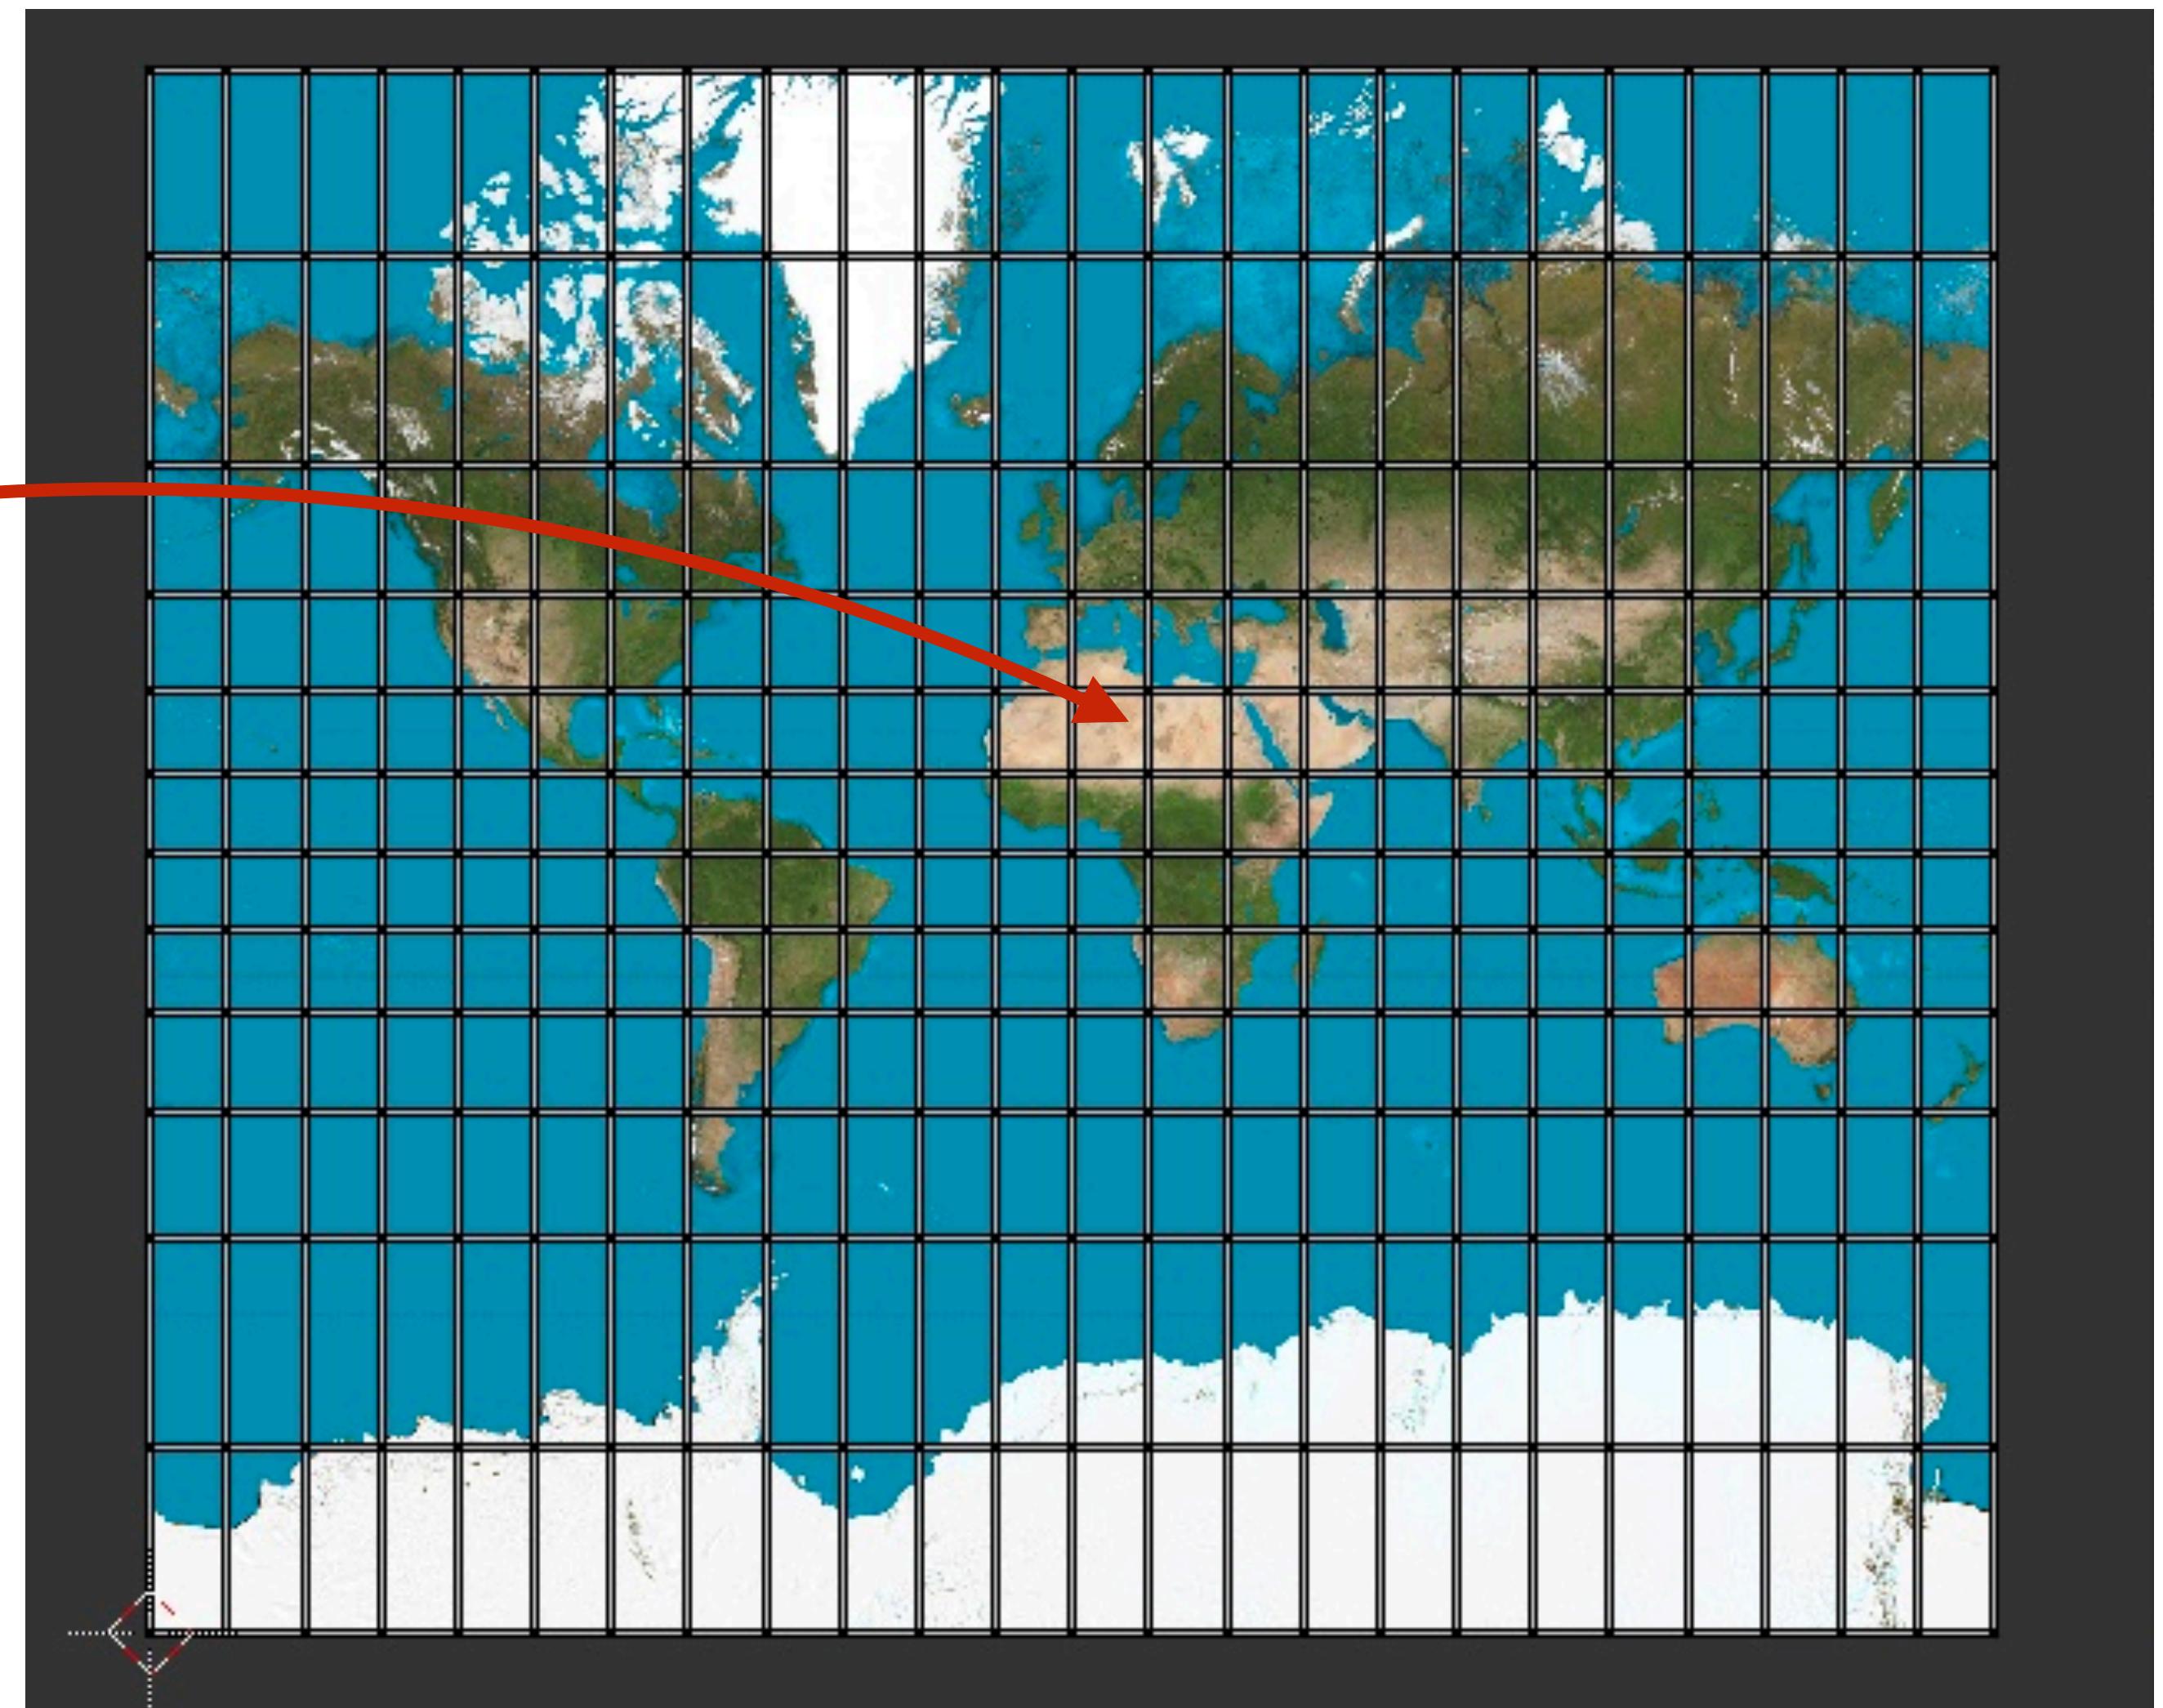
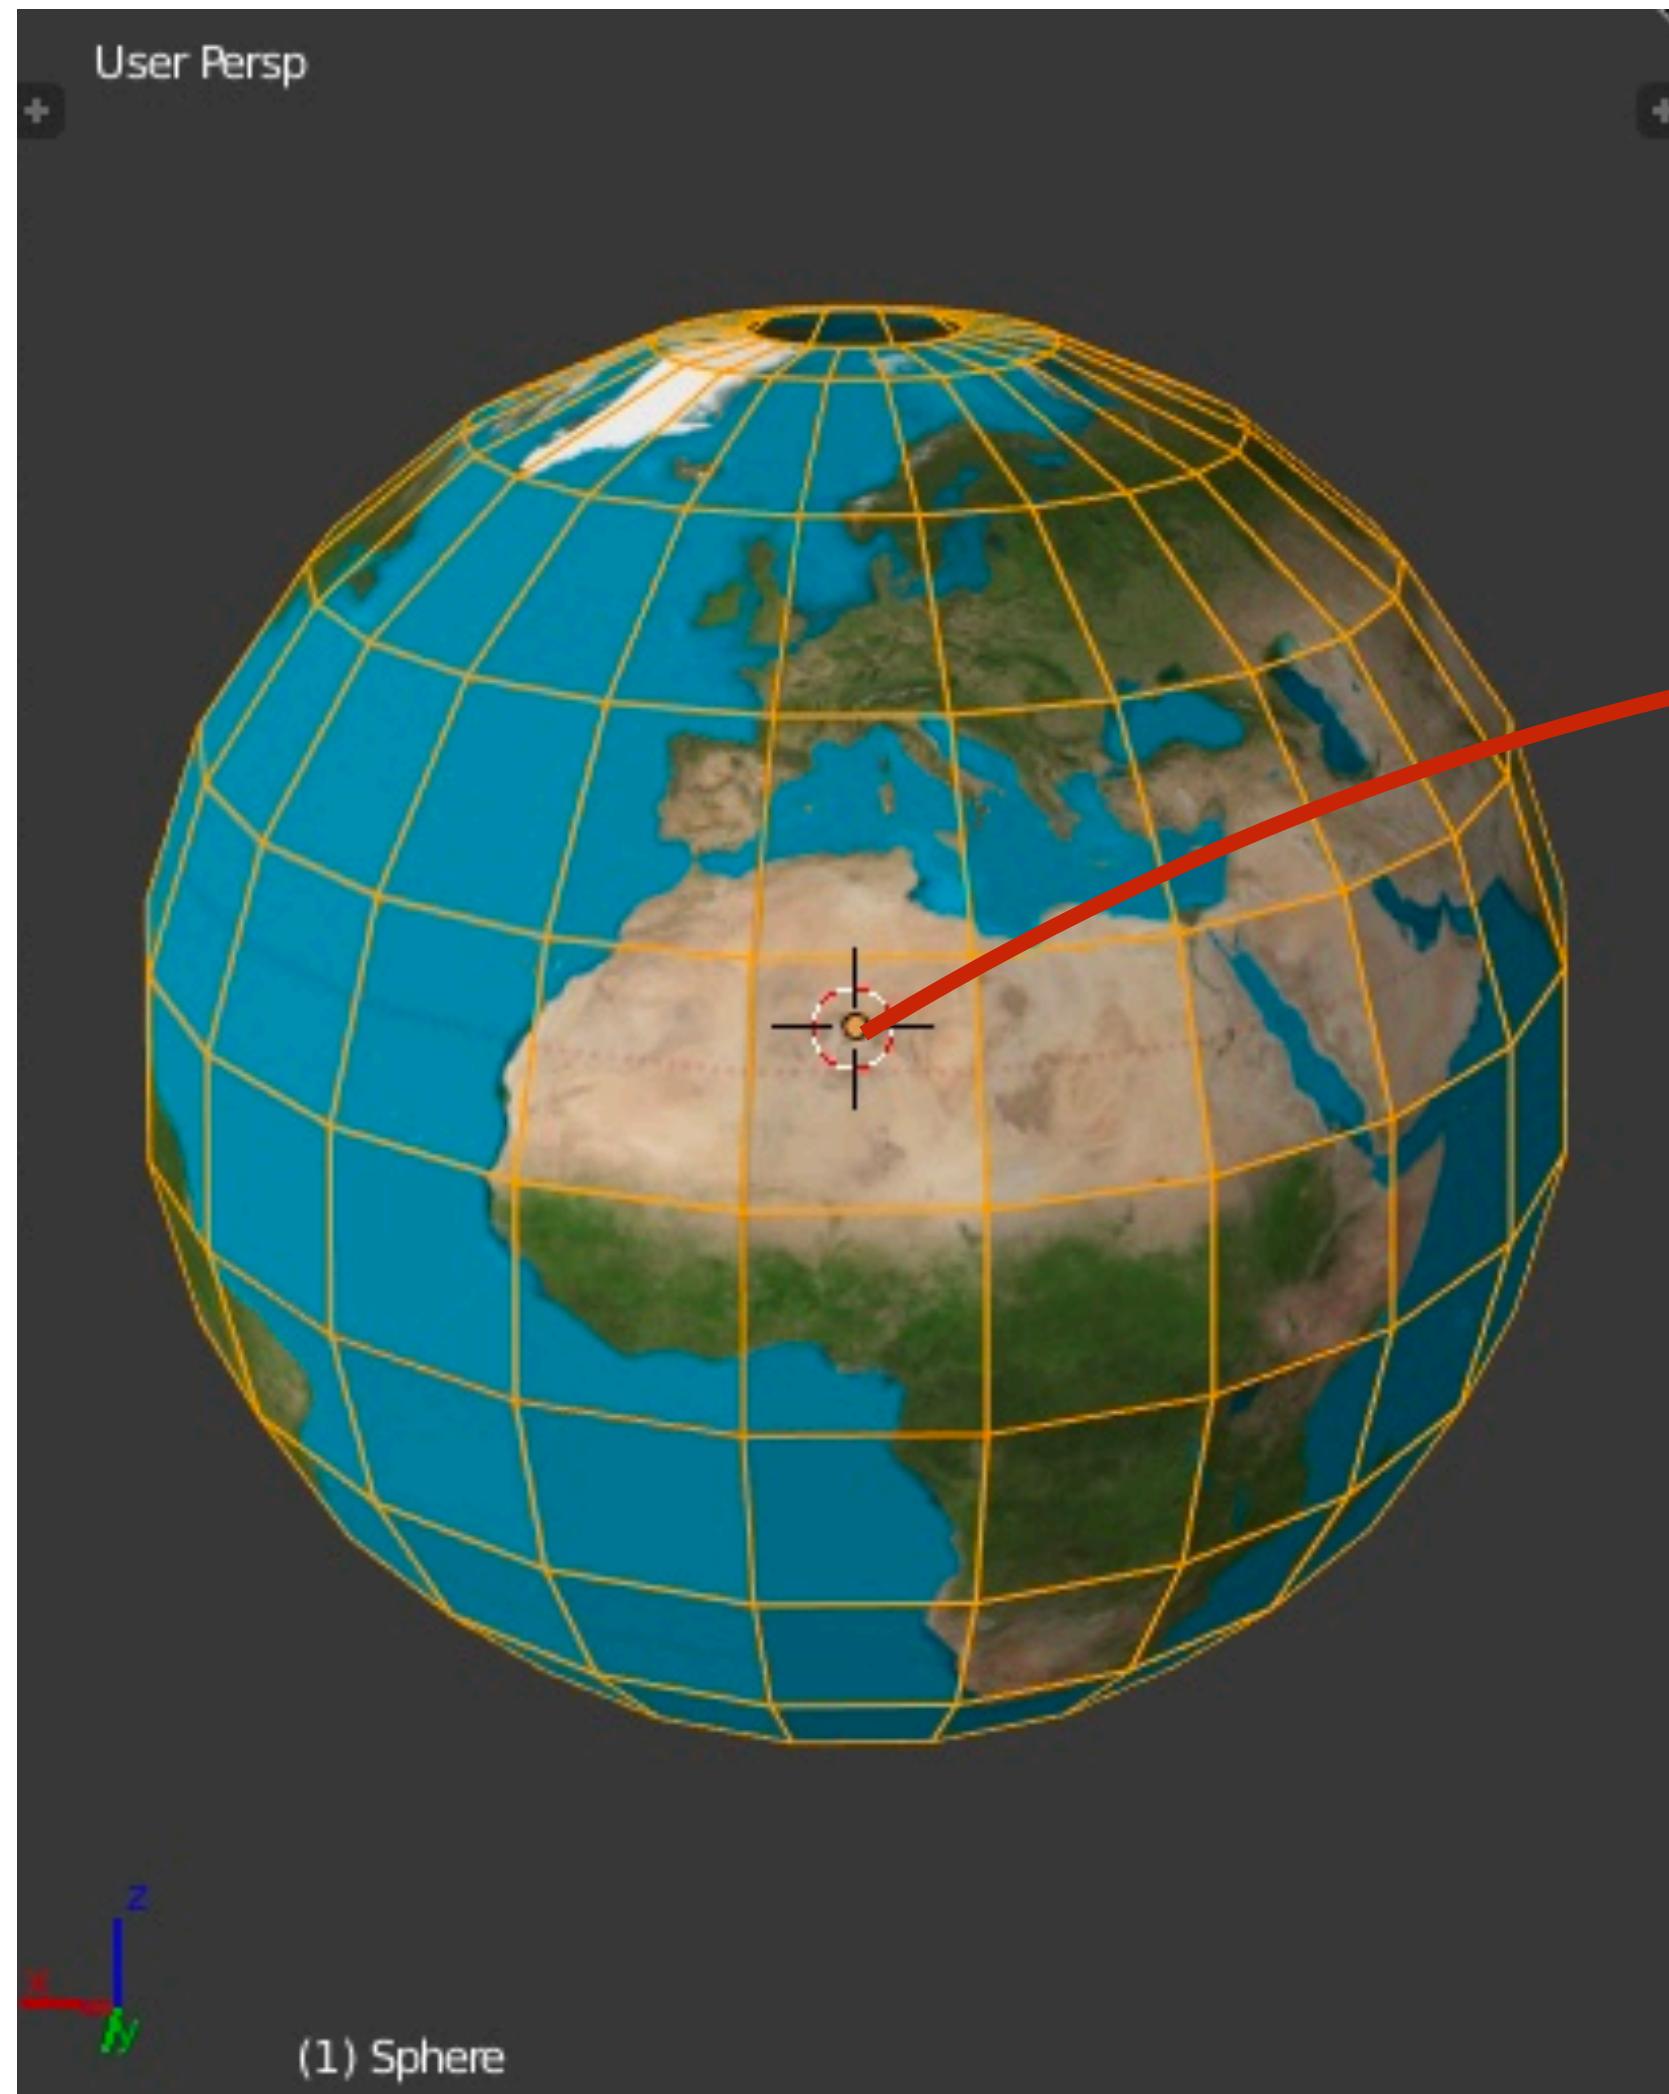


Final rendered result (entire cube shown).

Location of triangle after projection onto screen shown in red.

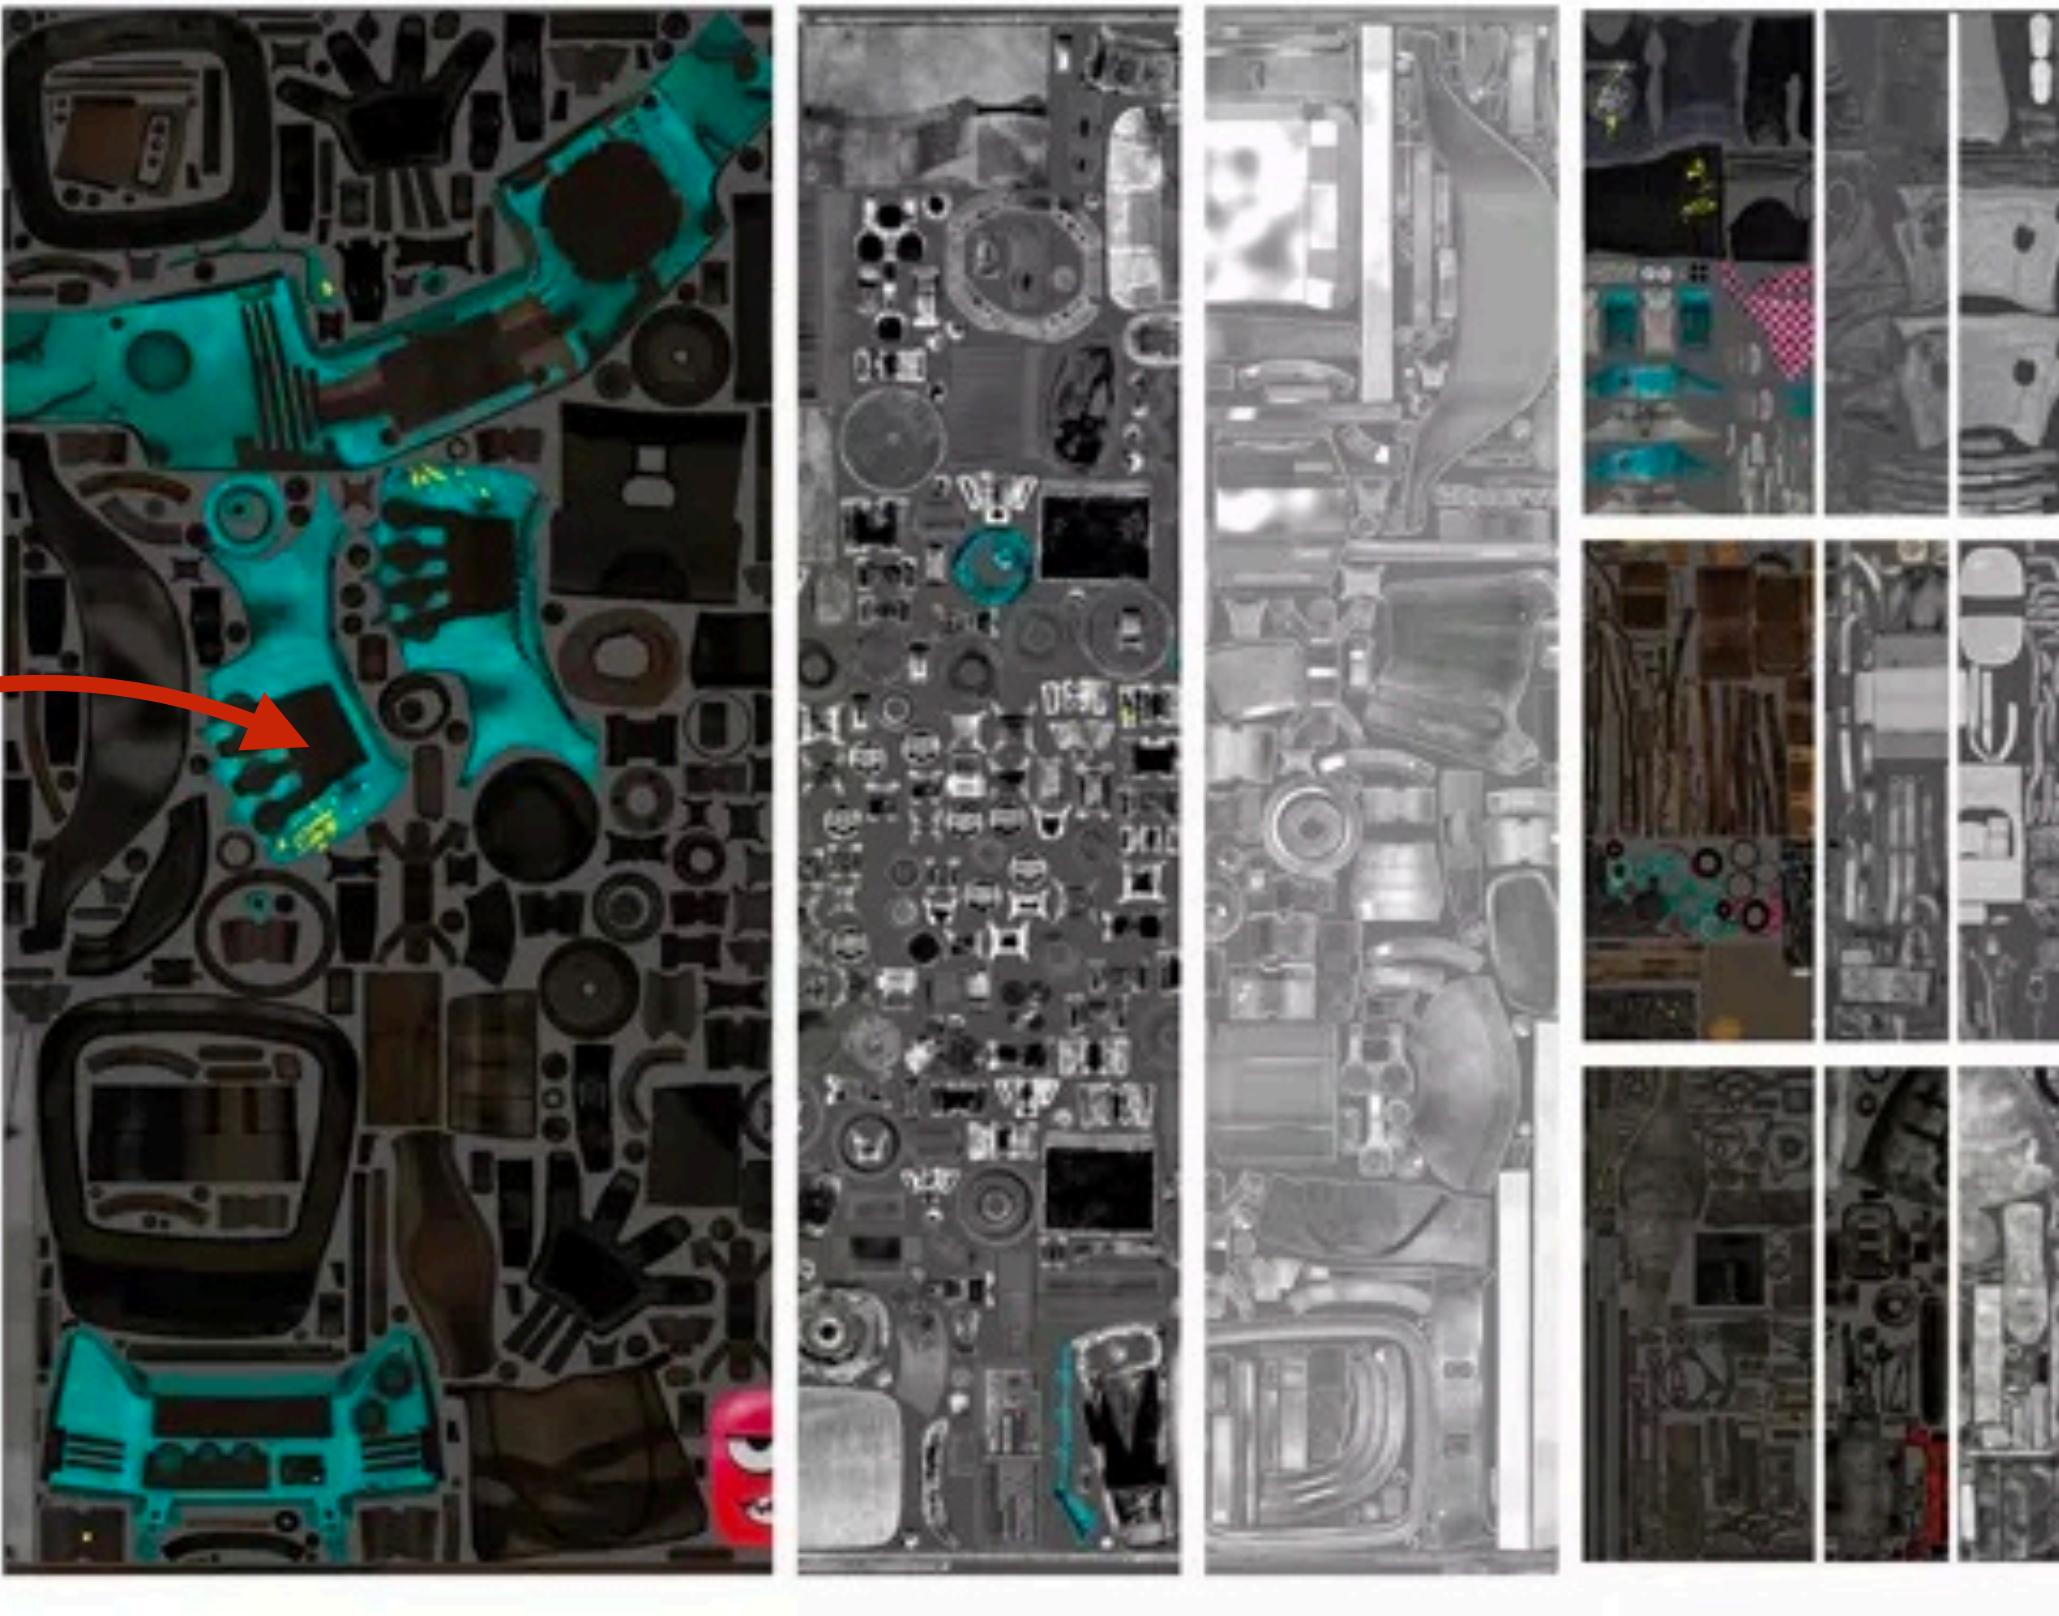
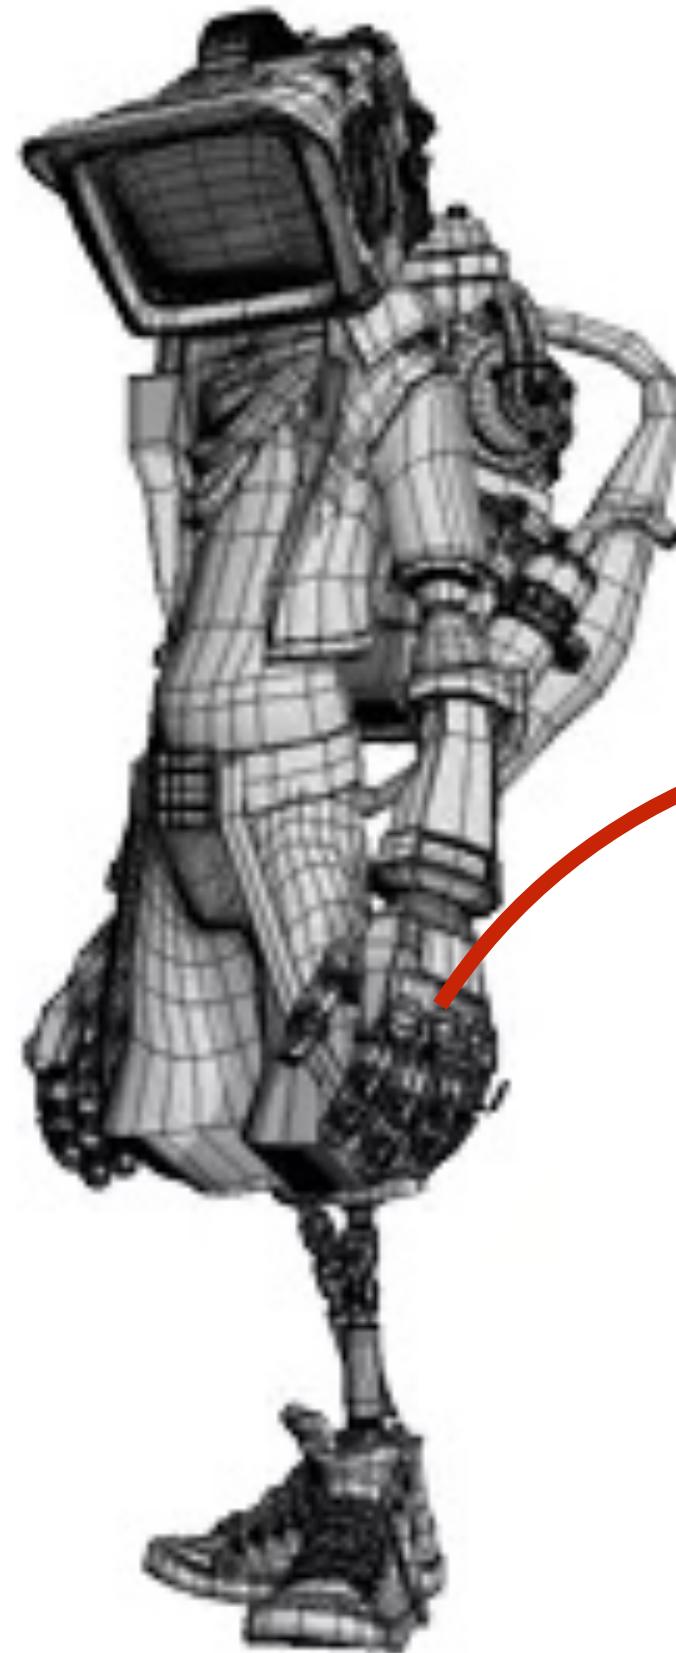
Today we'll assume surface-to-texture space mapping is provided as per vertex attribute
(Not discussing methods for generating surface texture parameterizations)

Many different mappings of surface position to texture space



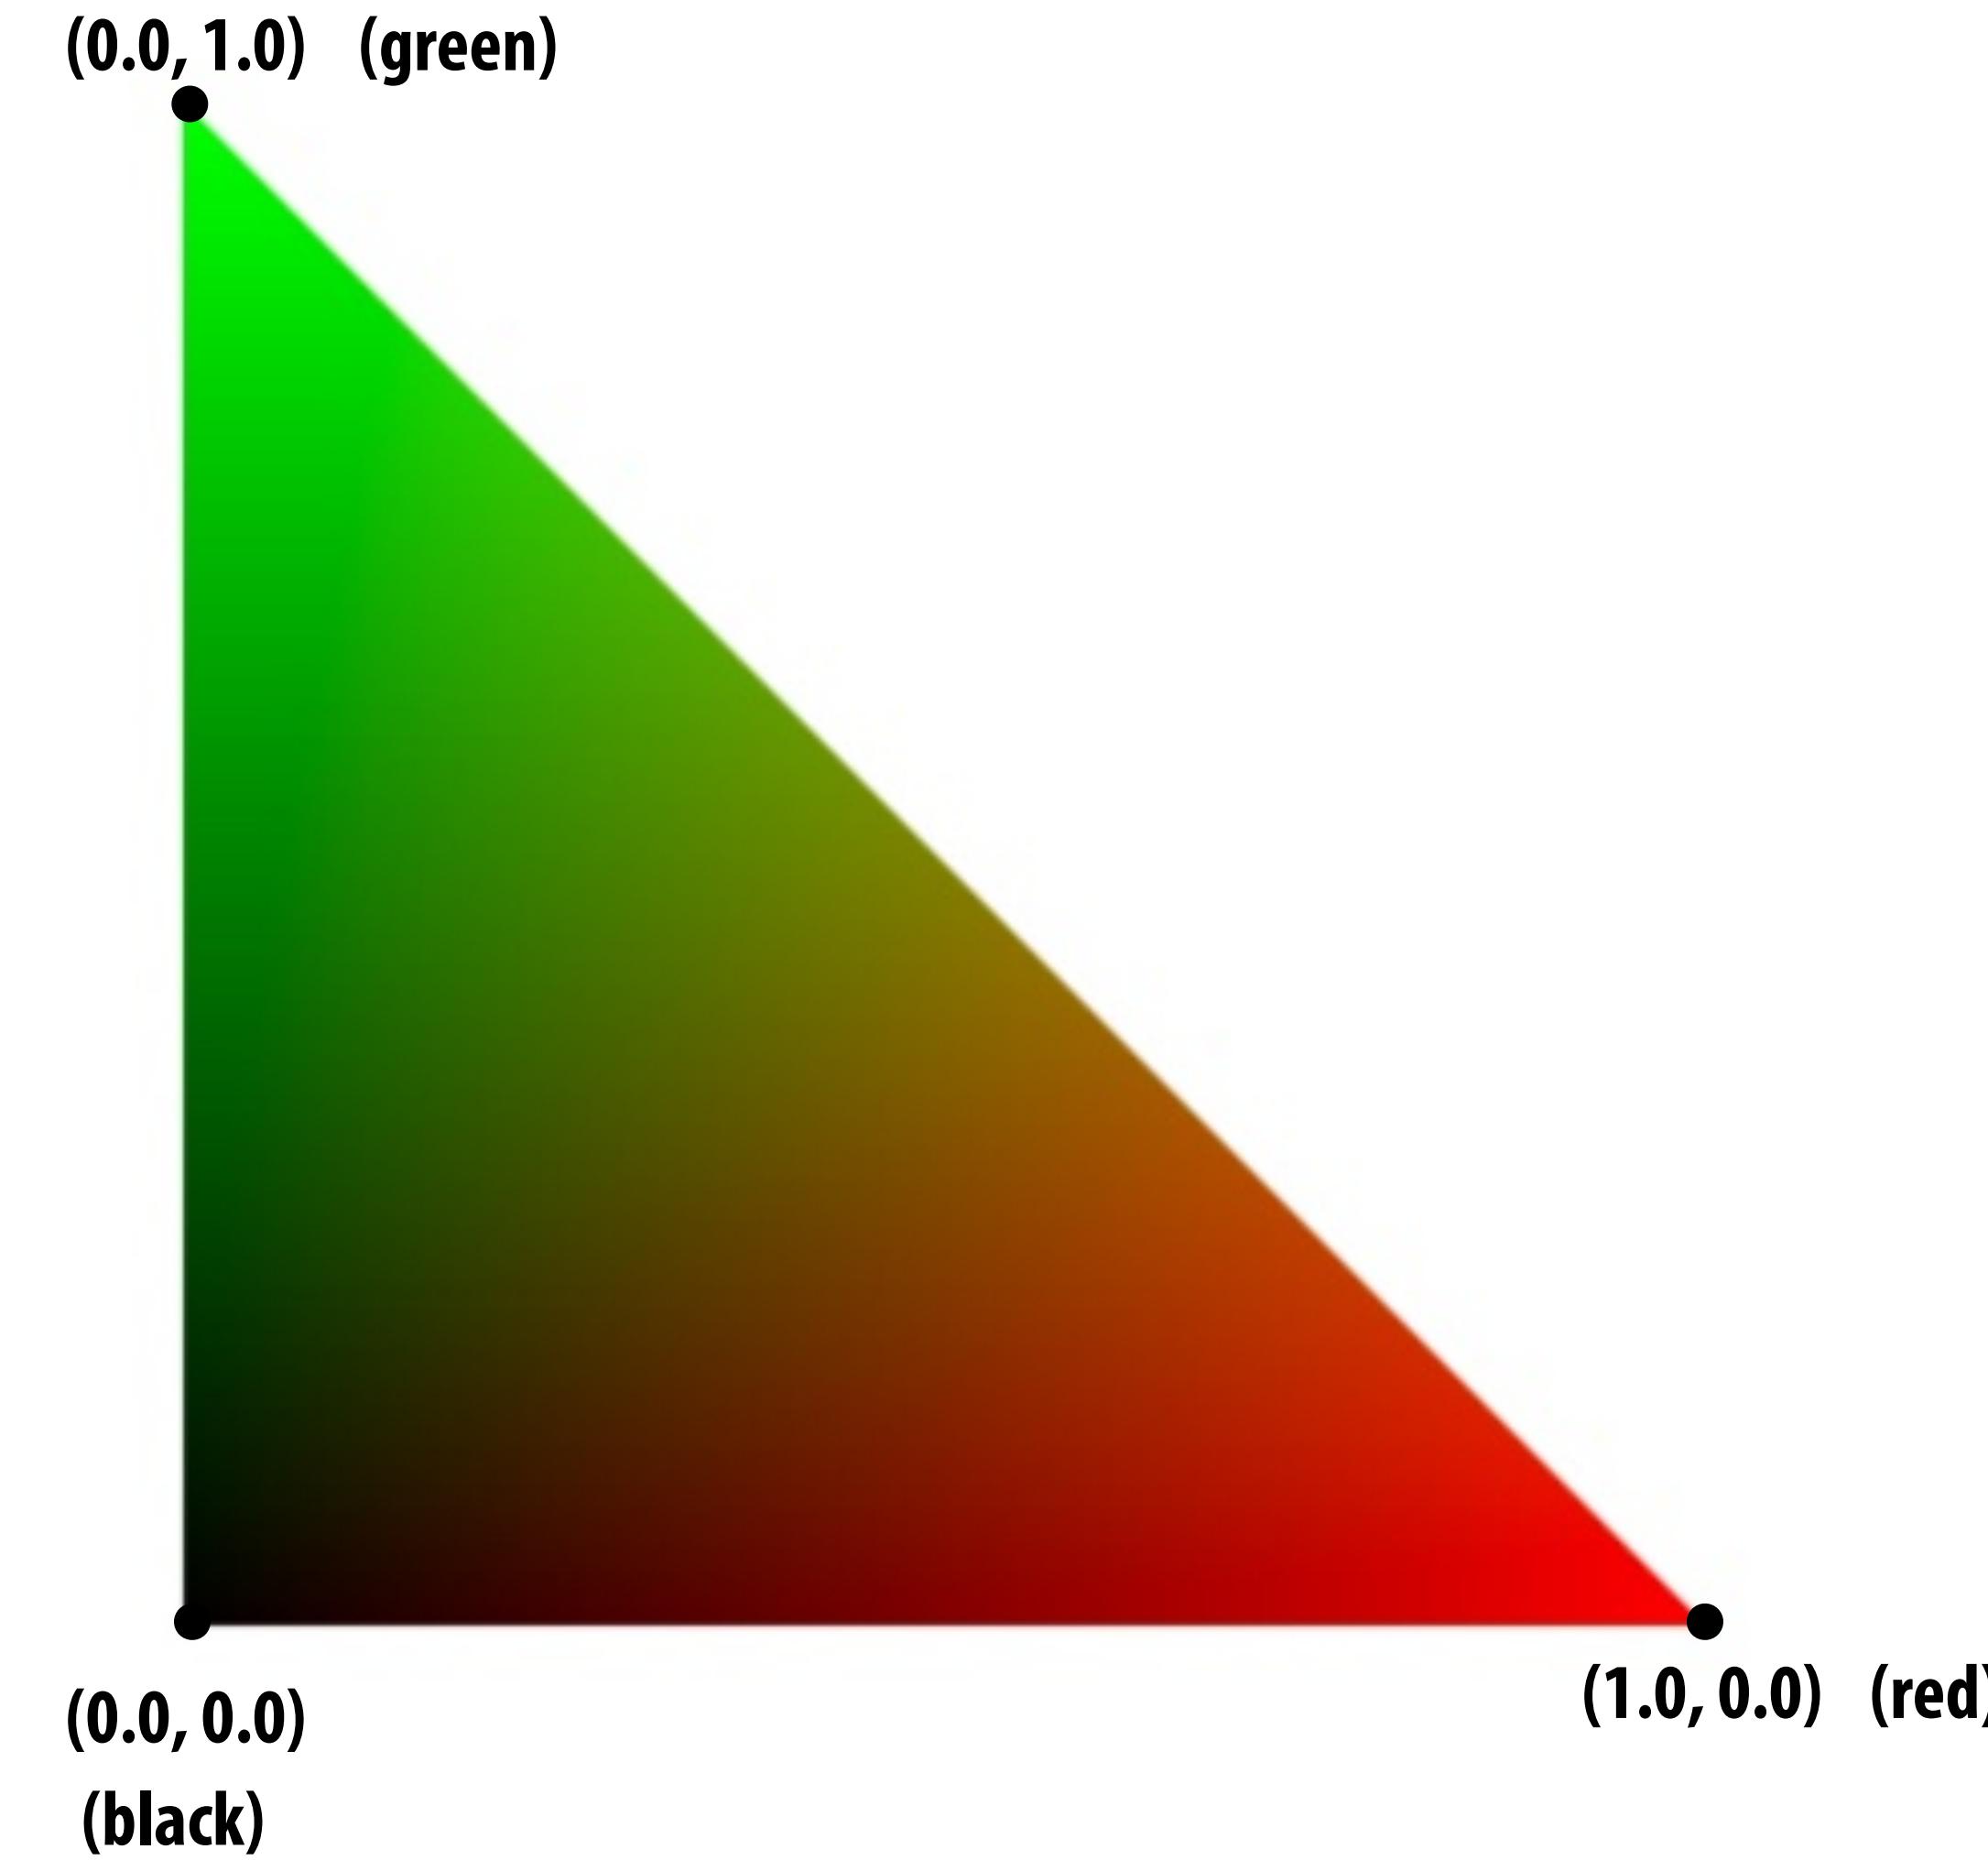
Example: mercator projection onto sphere

Texture “atlas”



Visualization of texture coordinates

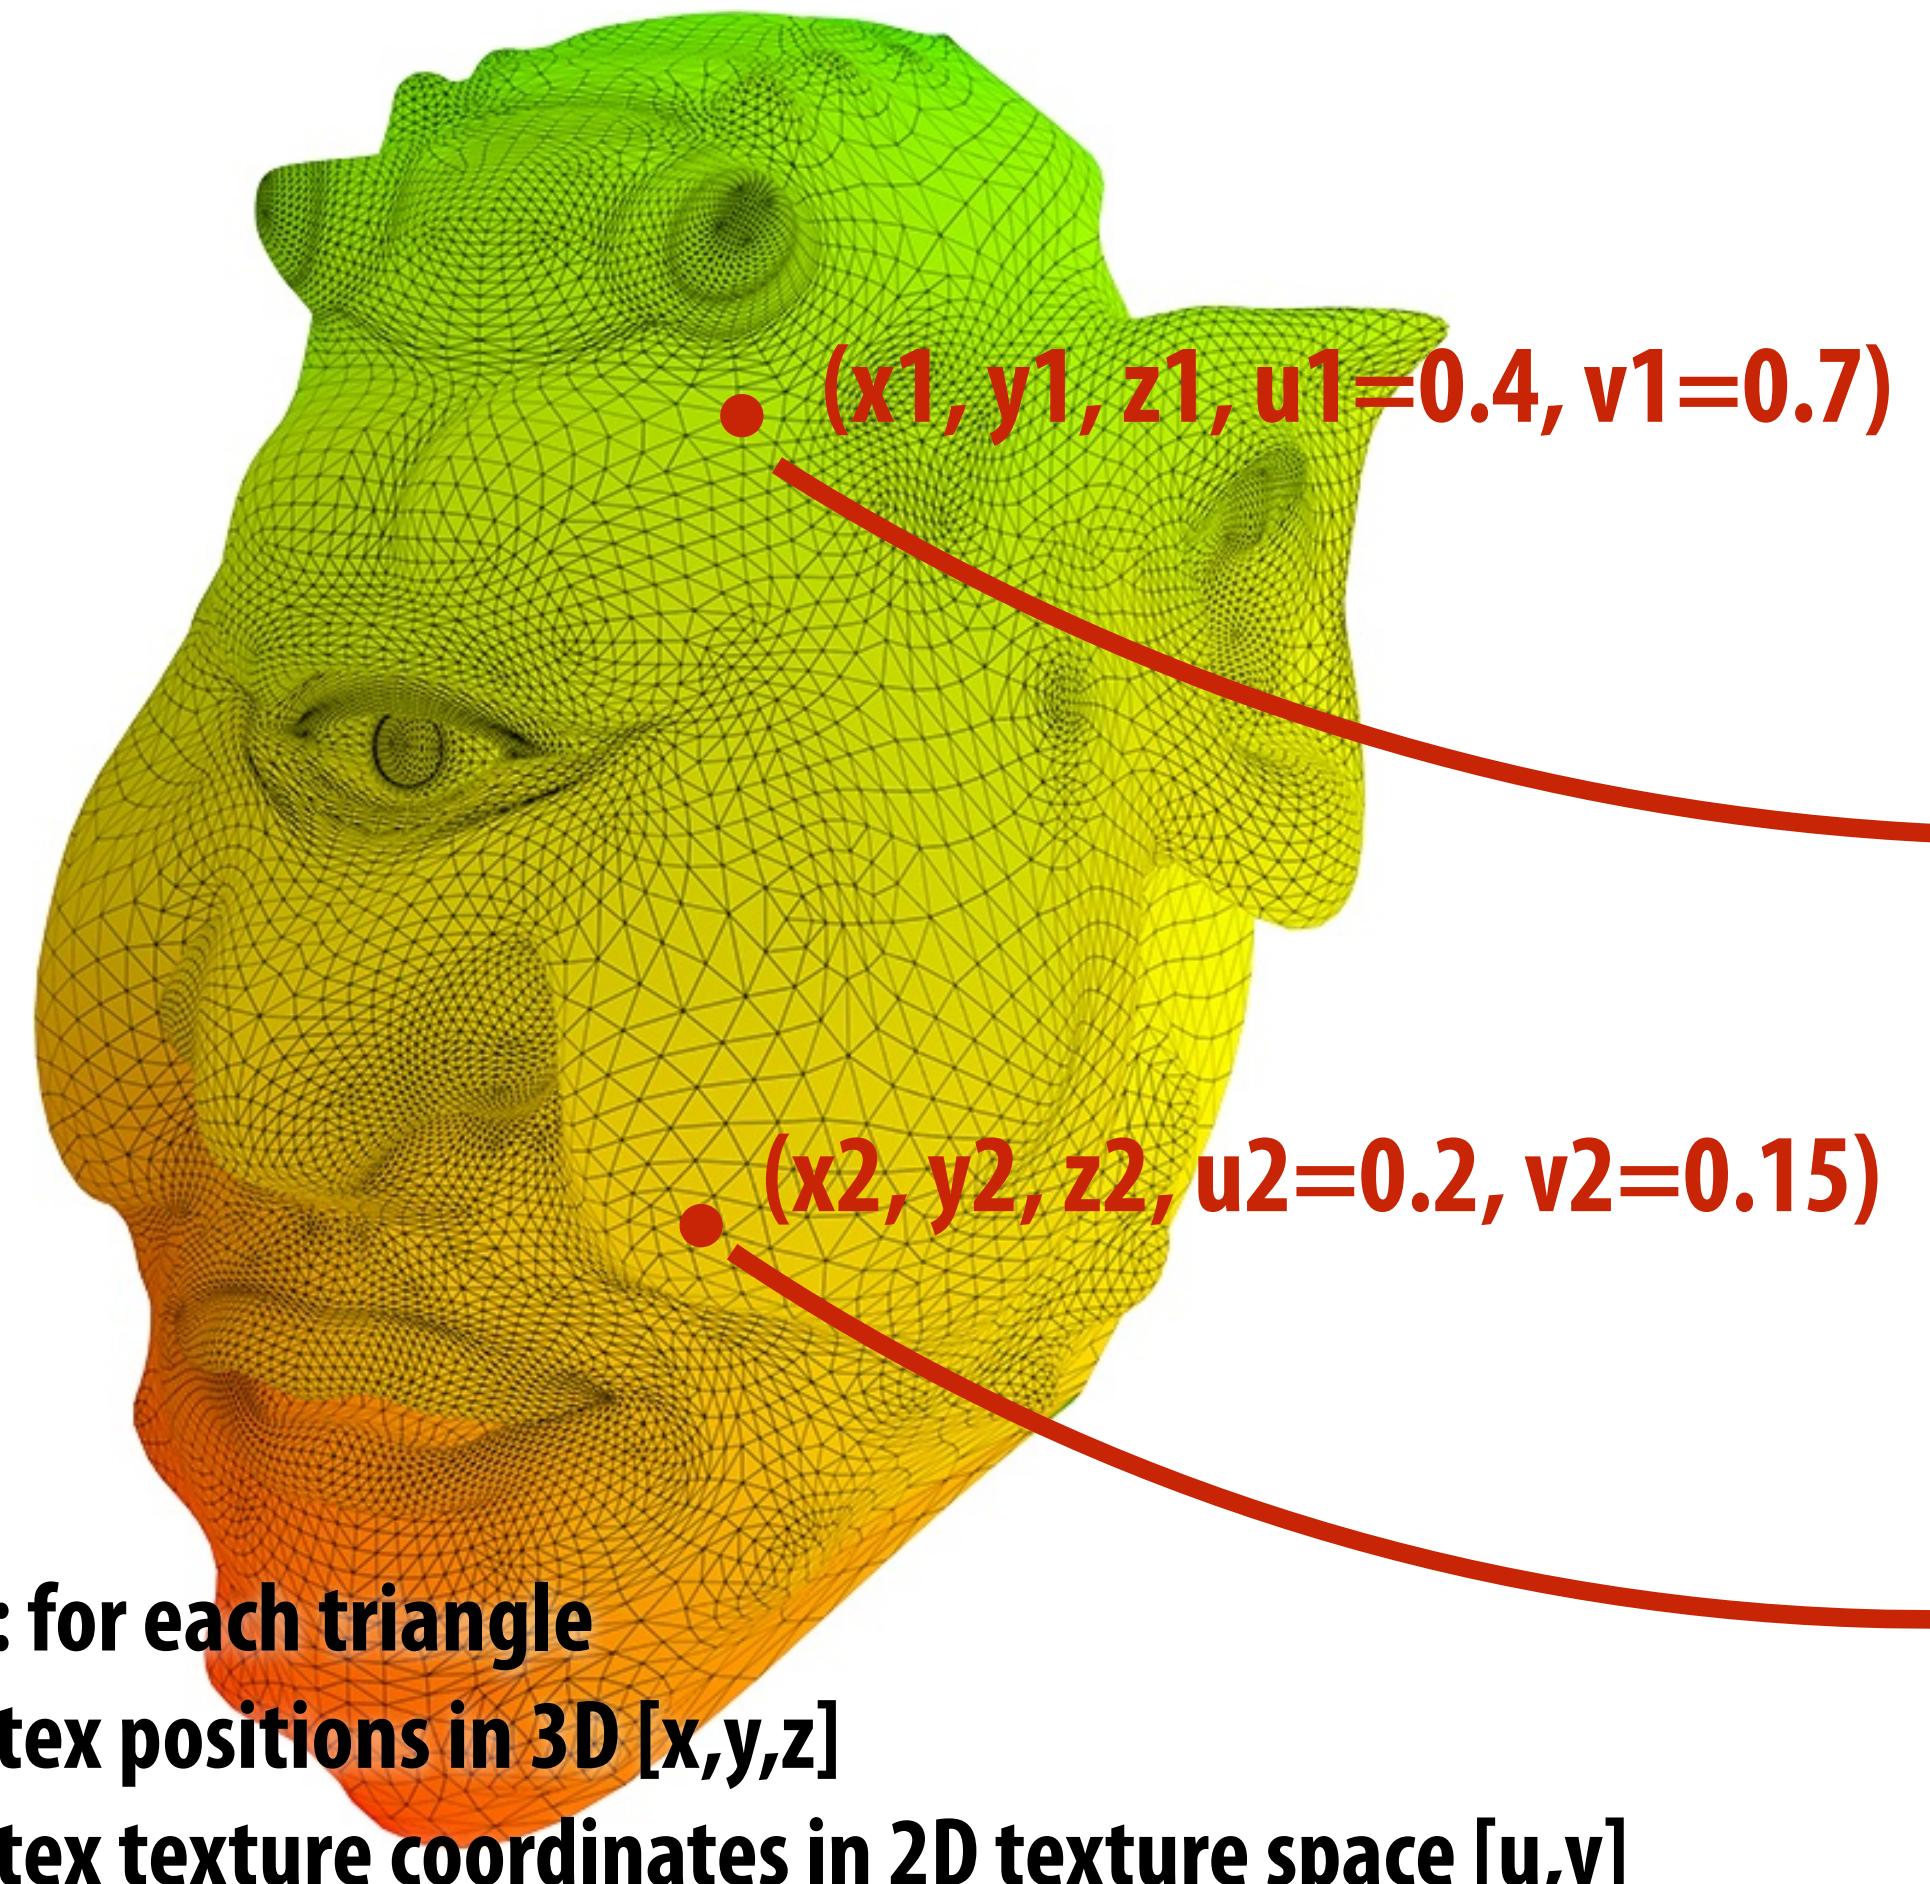
Texture coordinates linearly interpolated over triangle



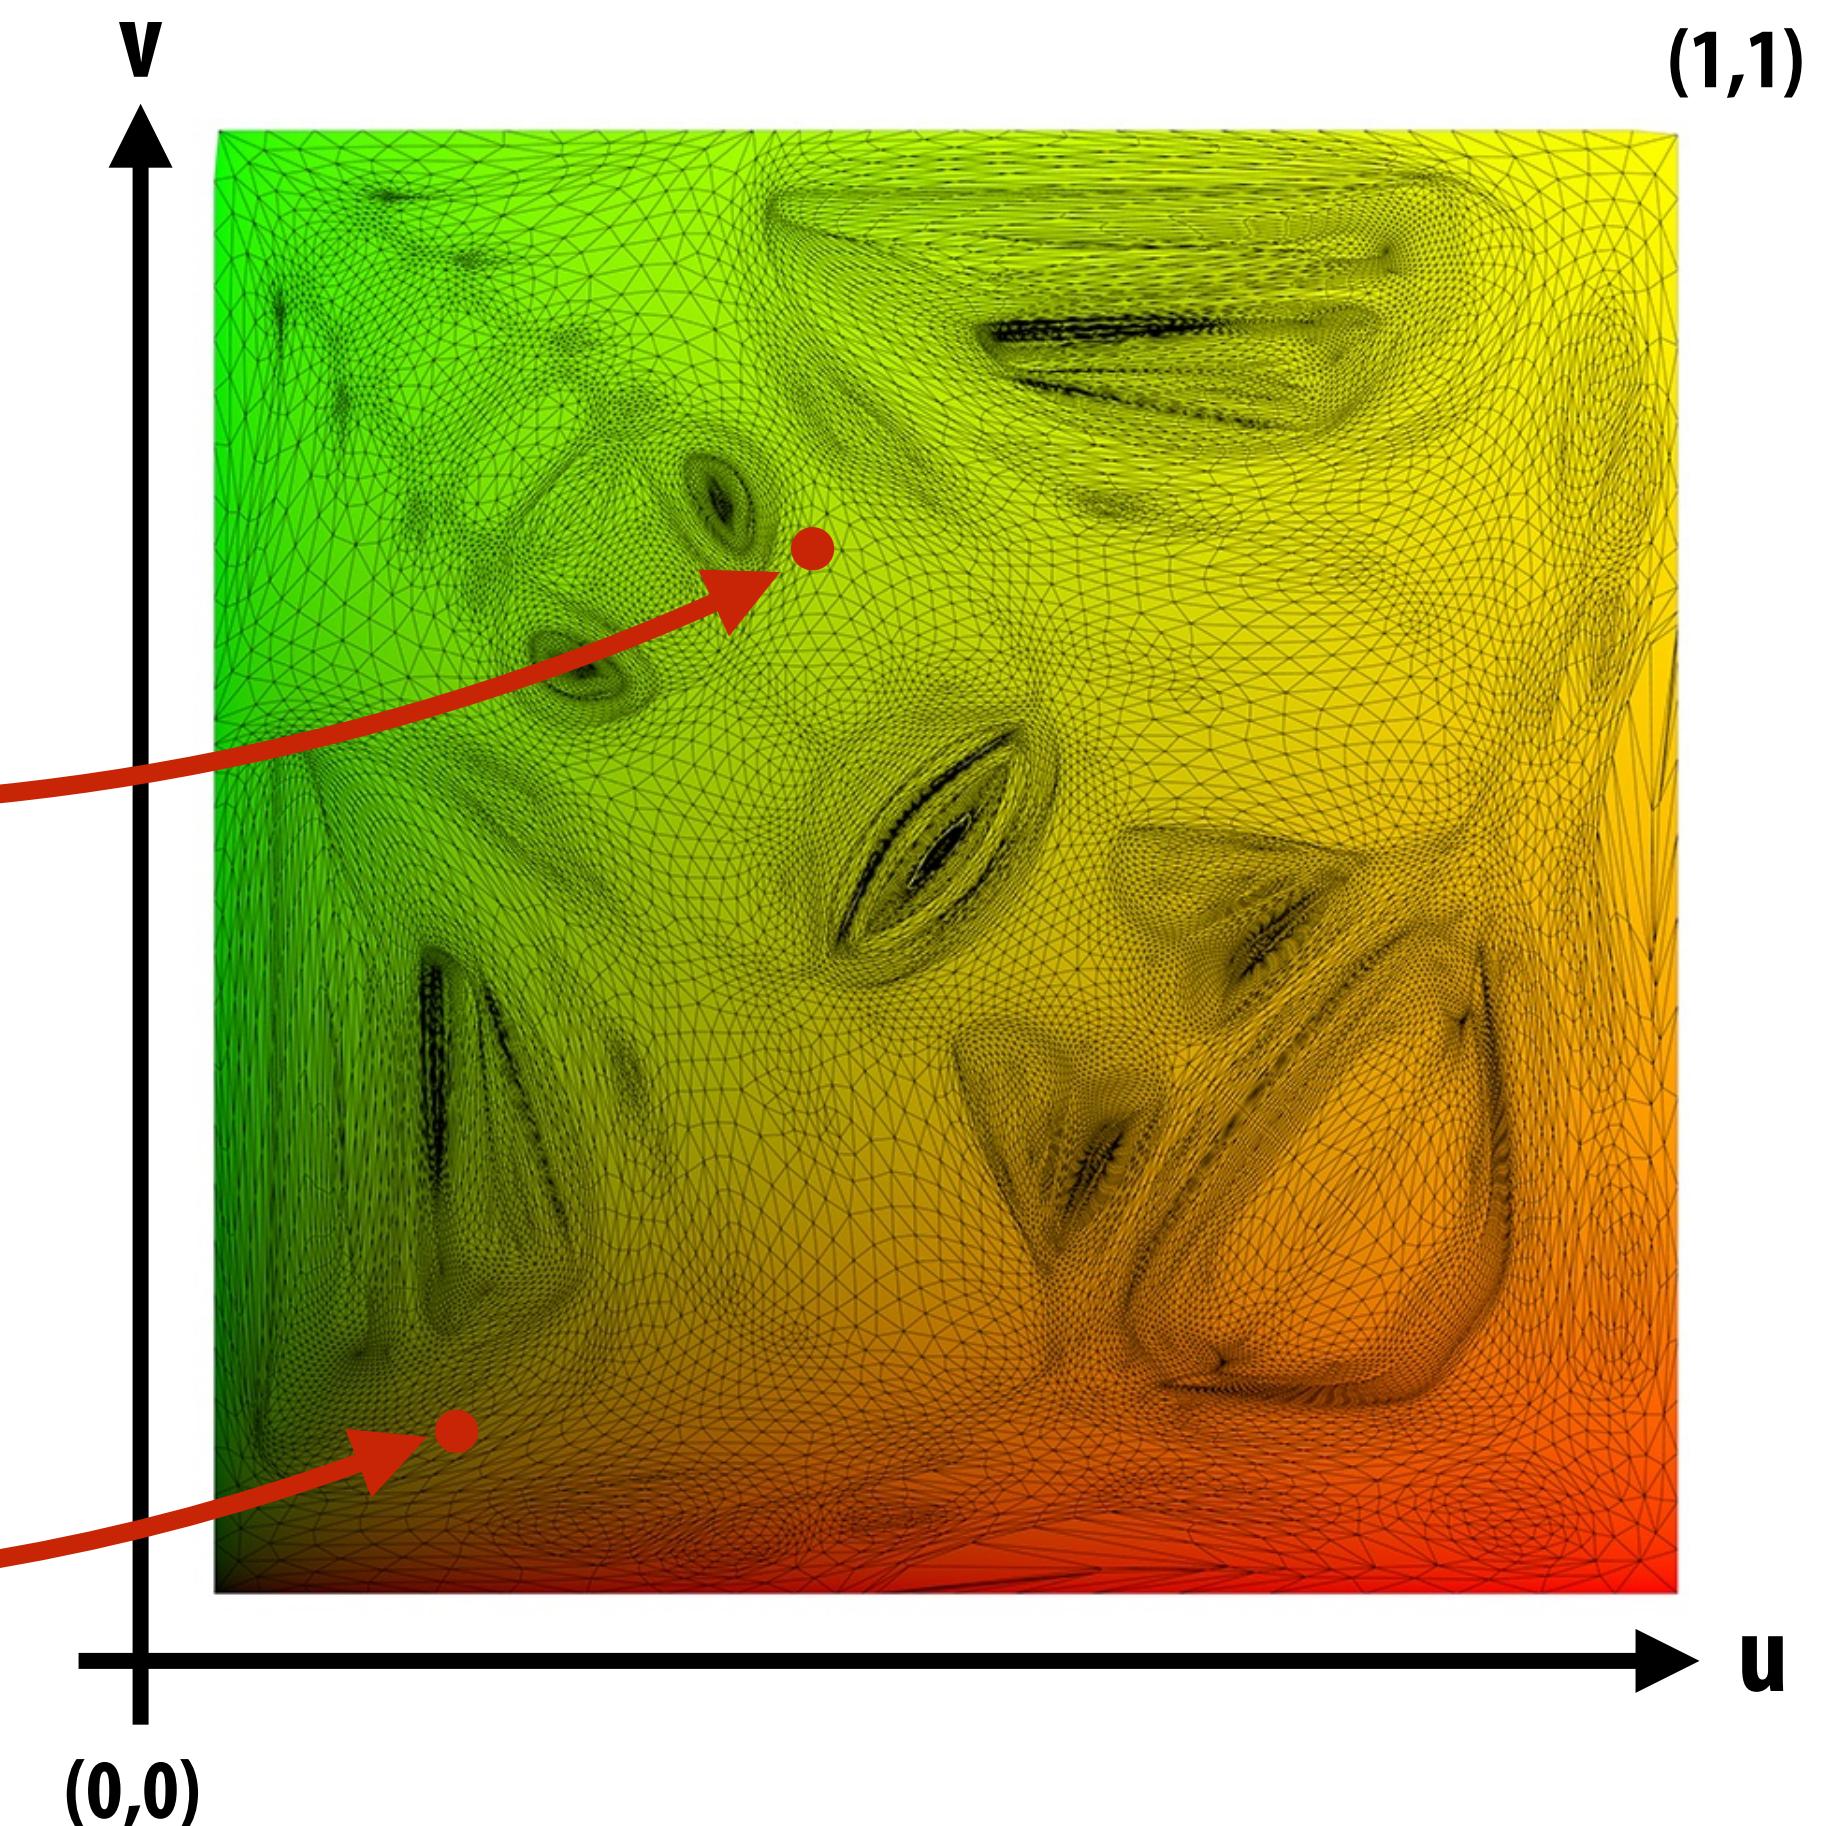
Texture coordinate values provided at triangle vertices

(Just like 3D positions are provided at vertices)

Visualization of texture coordinate value on mesh
(texture coordinate = color)

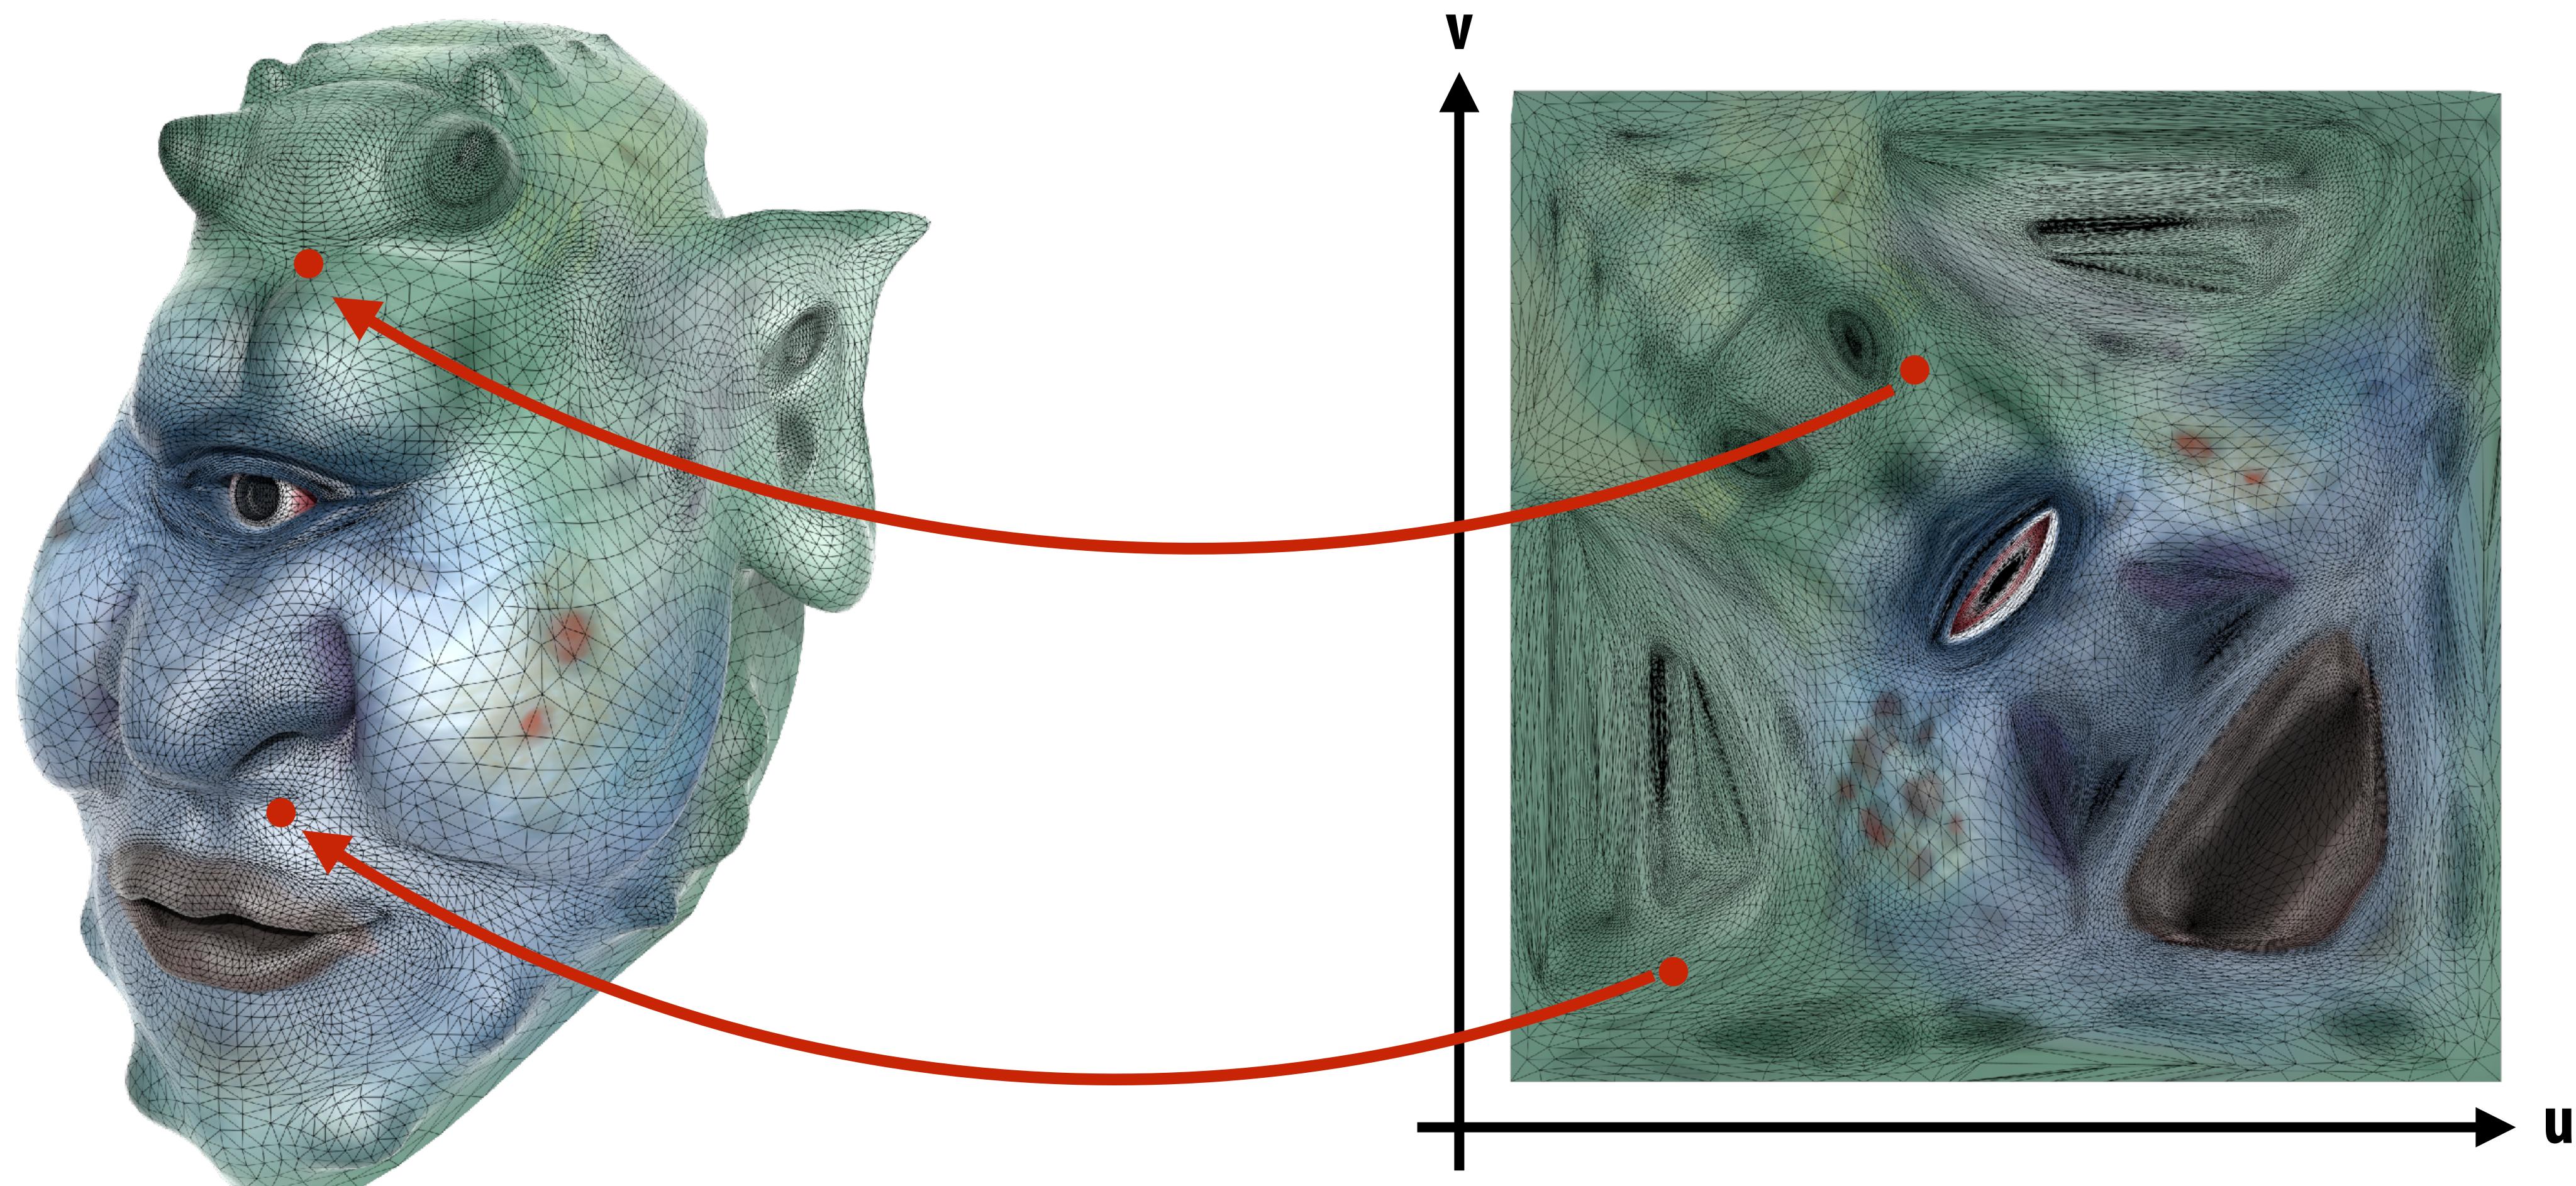


Visualization of location of triangle vertices
in texture space



Texture mapping adds detail

Sample texture map at specified location in *texture coordinate space* to determine the surface's color at the corresponding point on surface.



Texture mapping adds detail

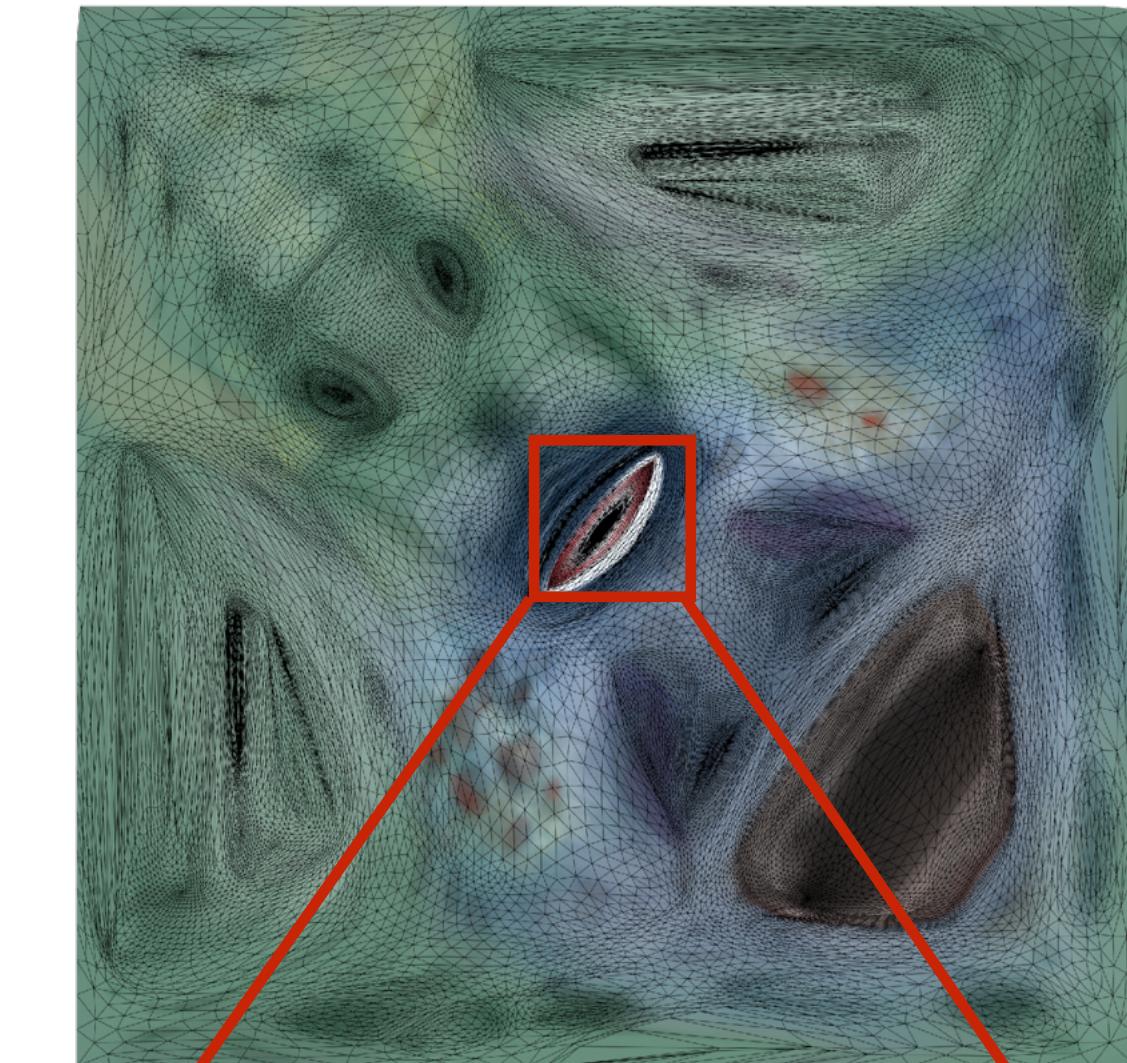
rendering without texture



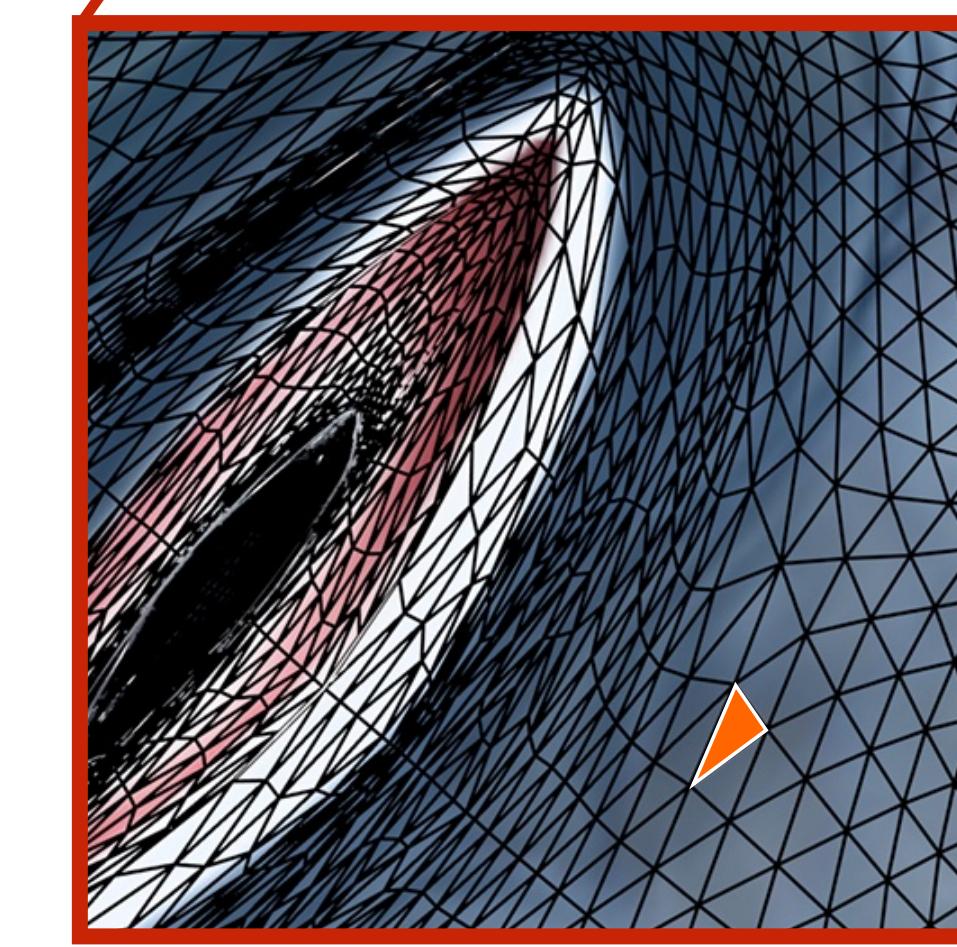
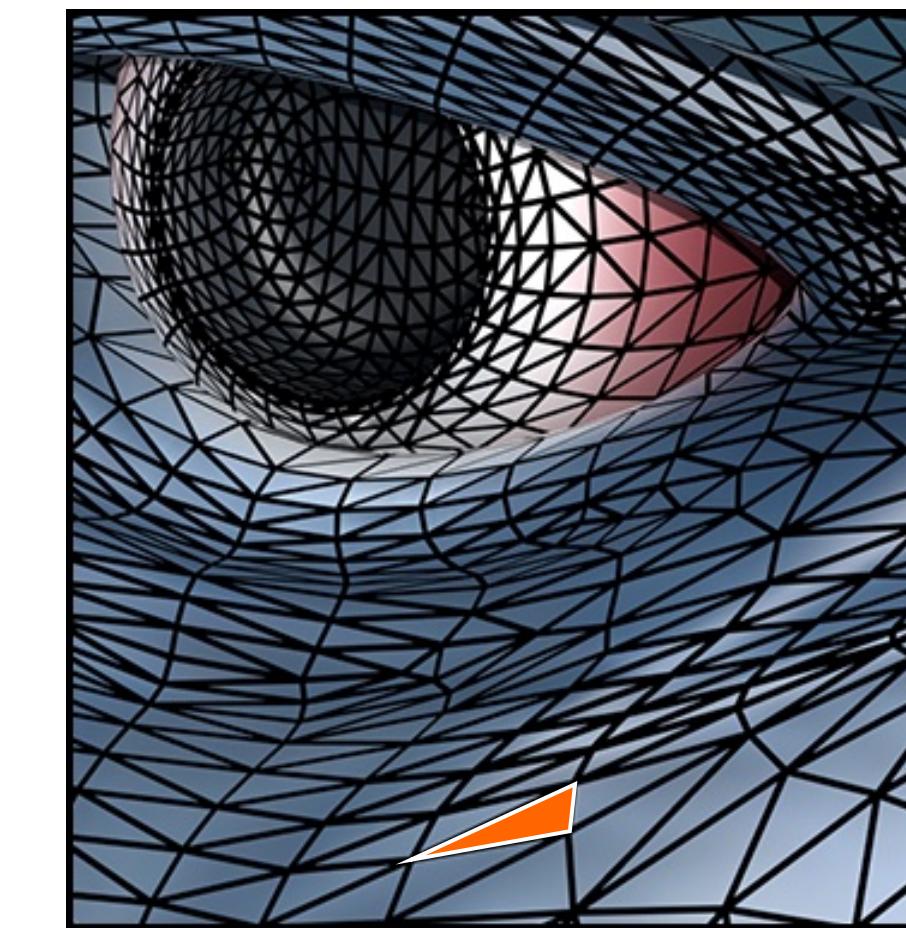
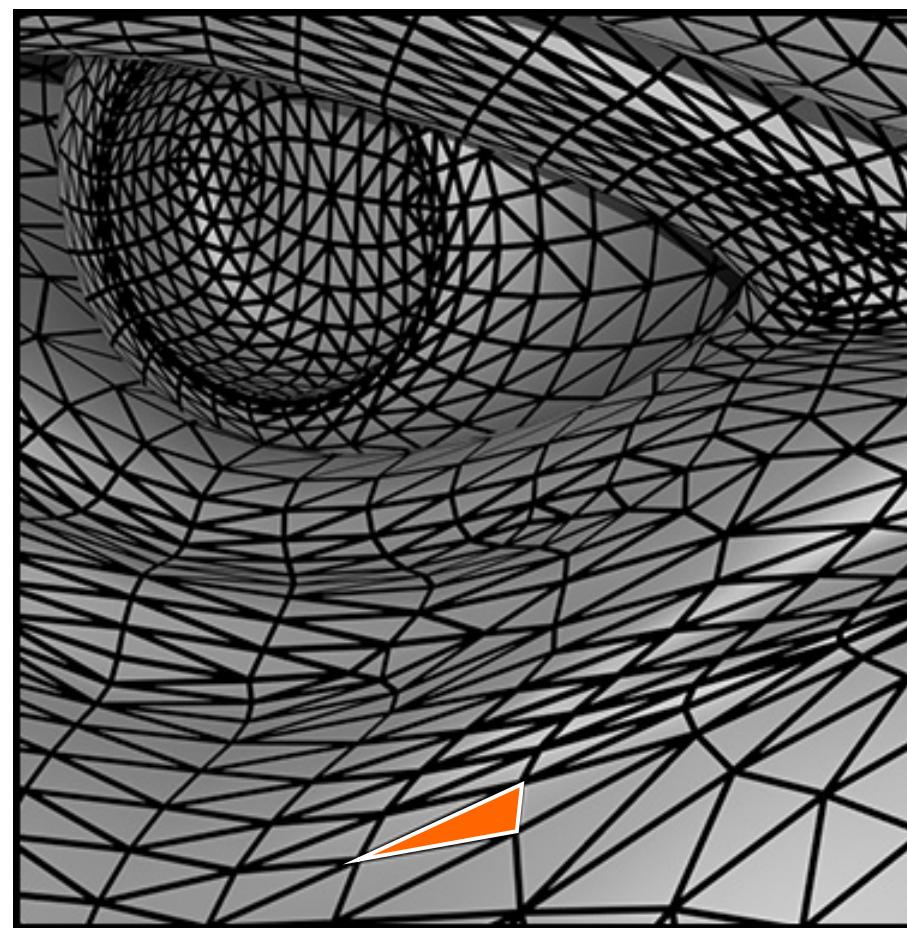
rendering with texture



texture image



zoom

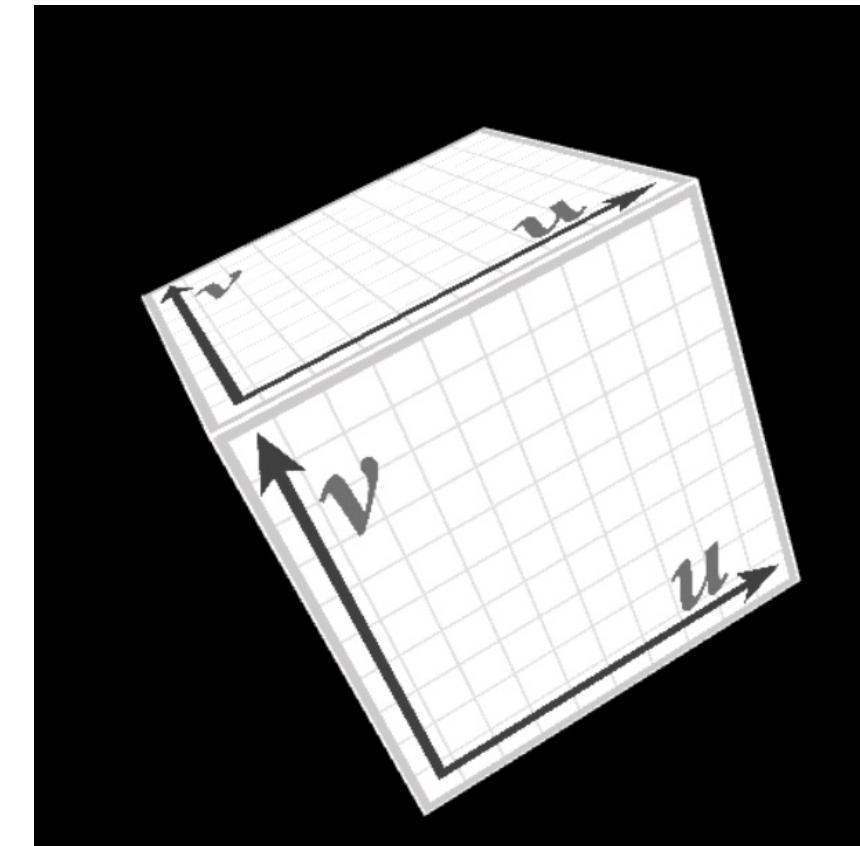
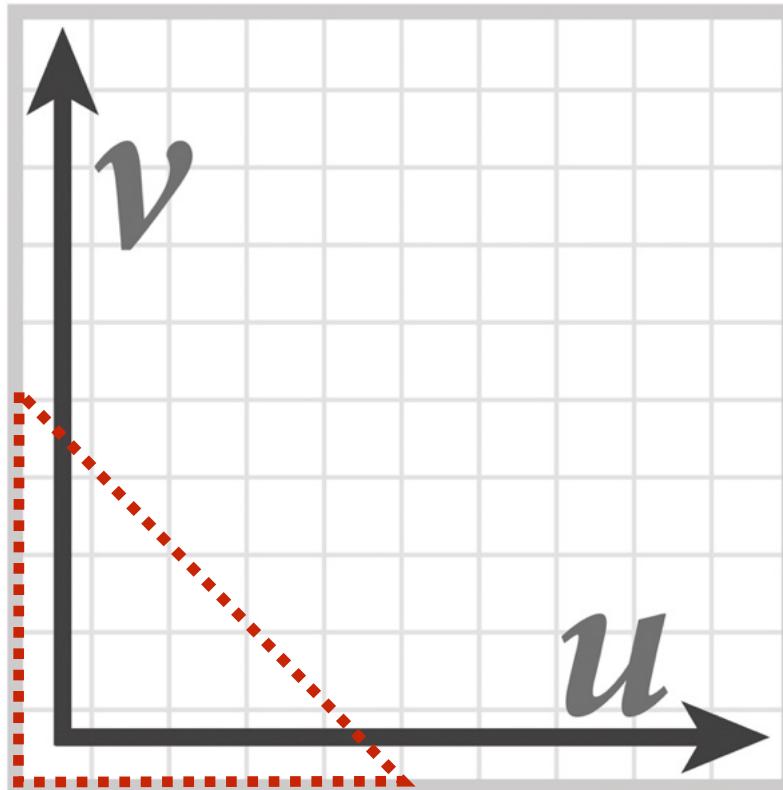
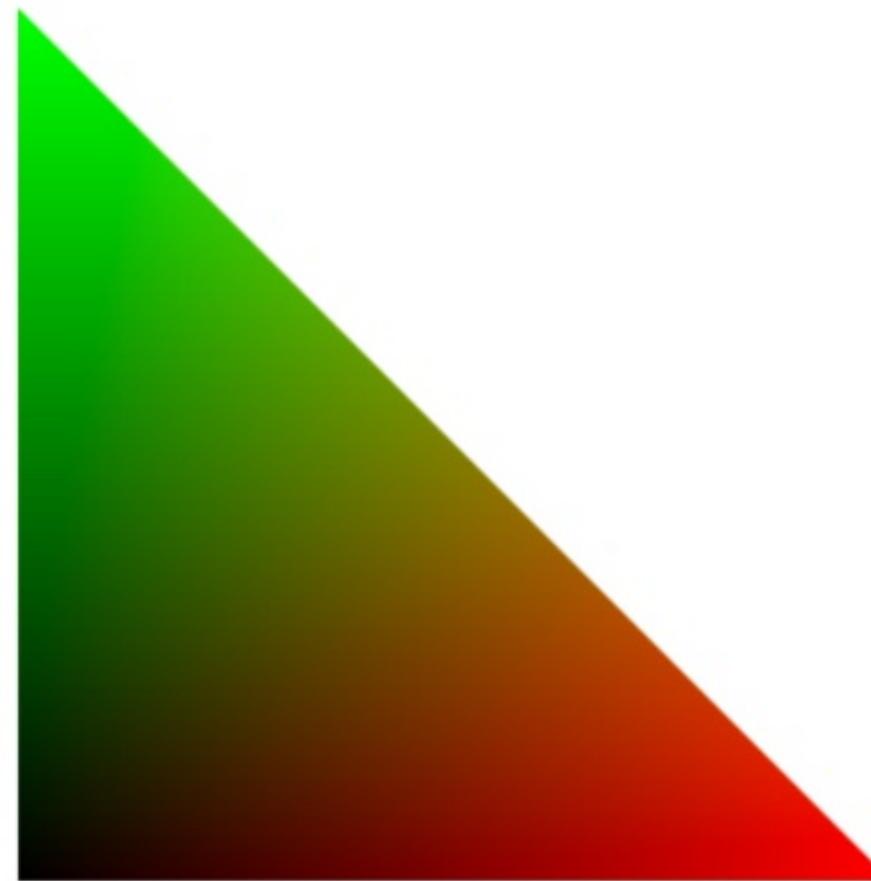


Each triangle “copies” a piece of the image back to the surface.

Texture sampling 101

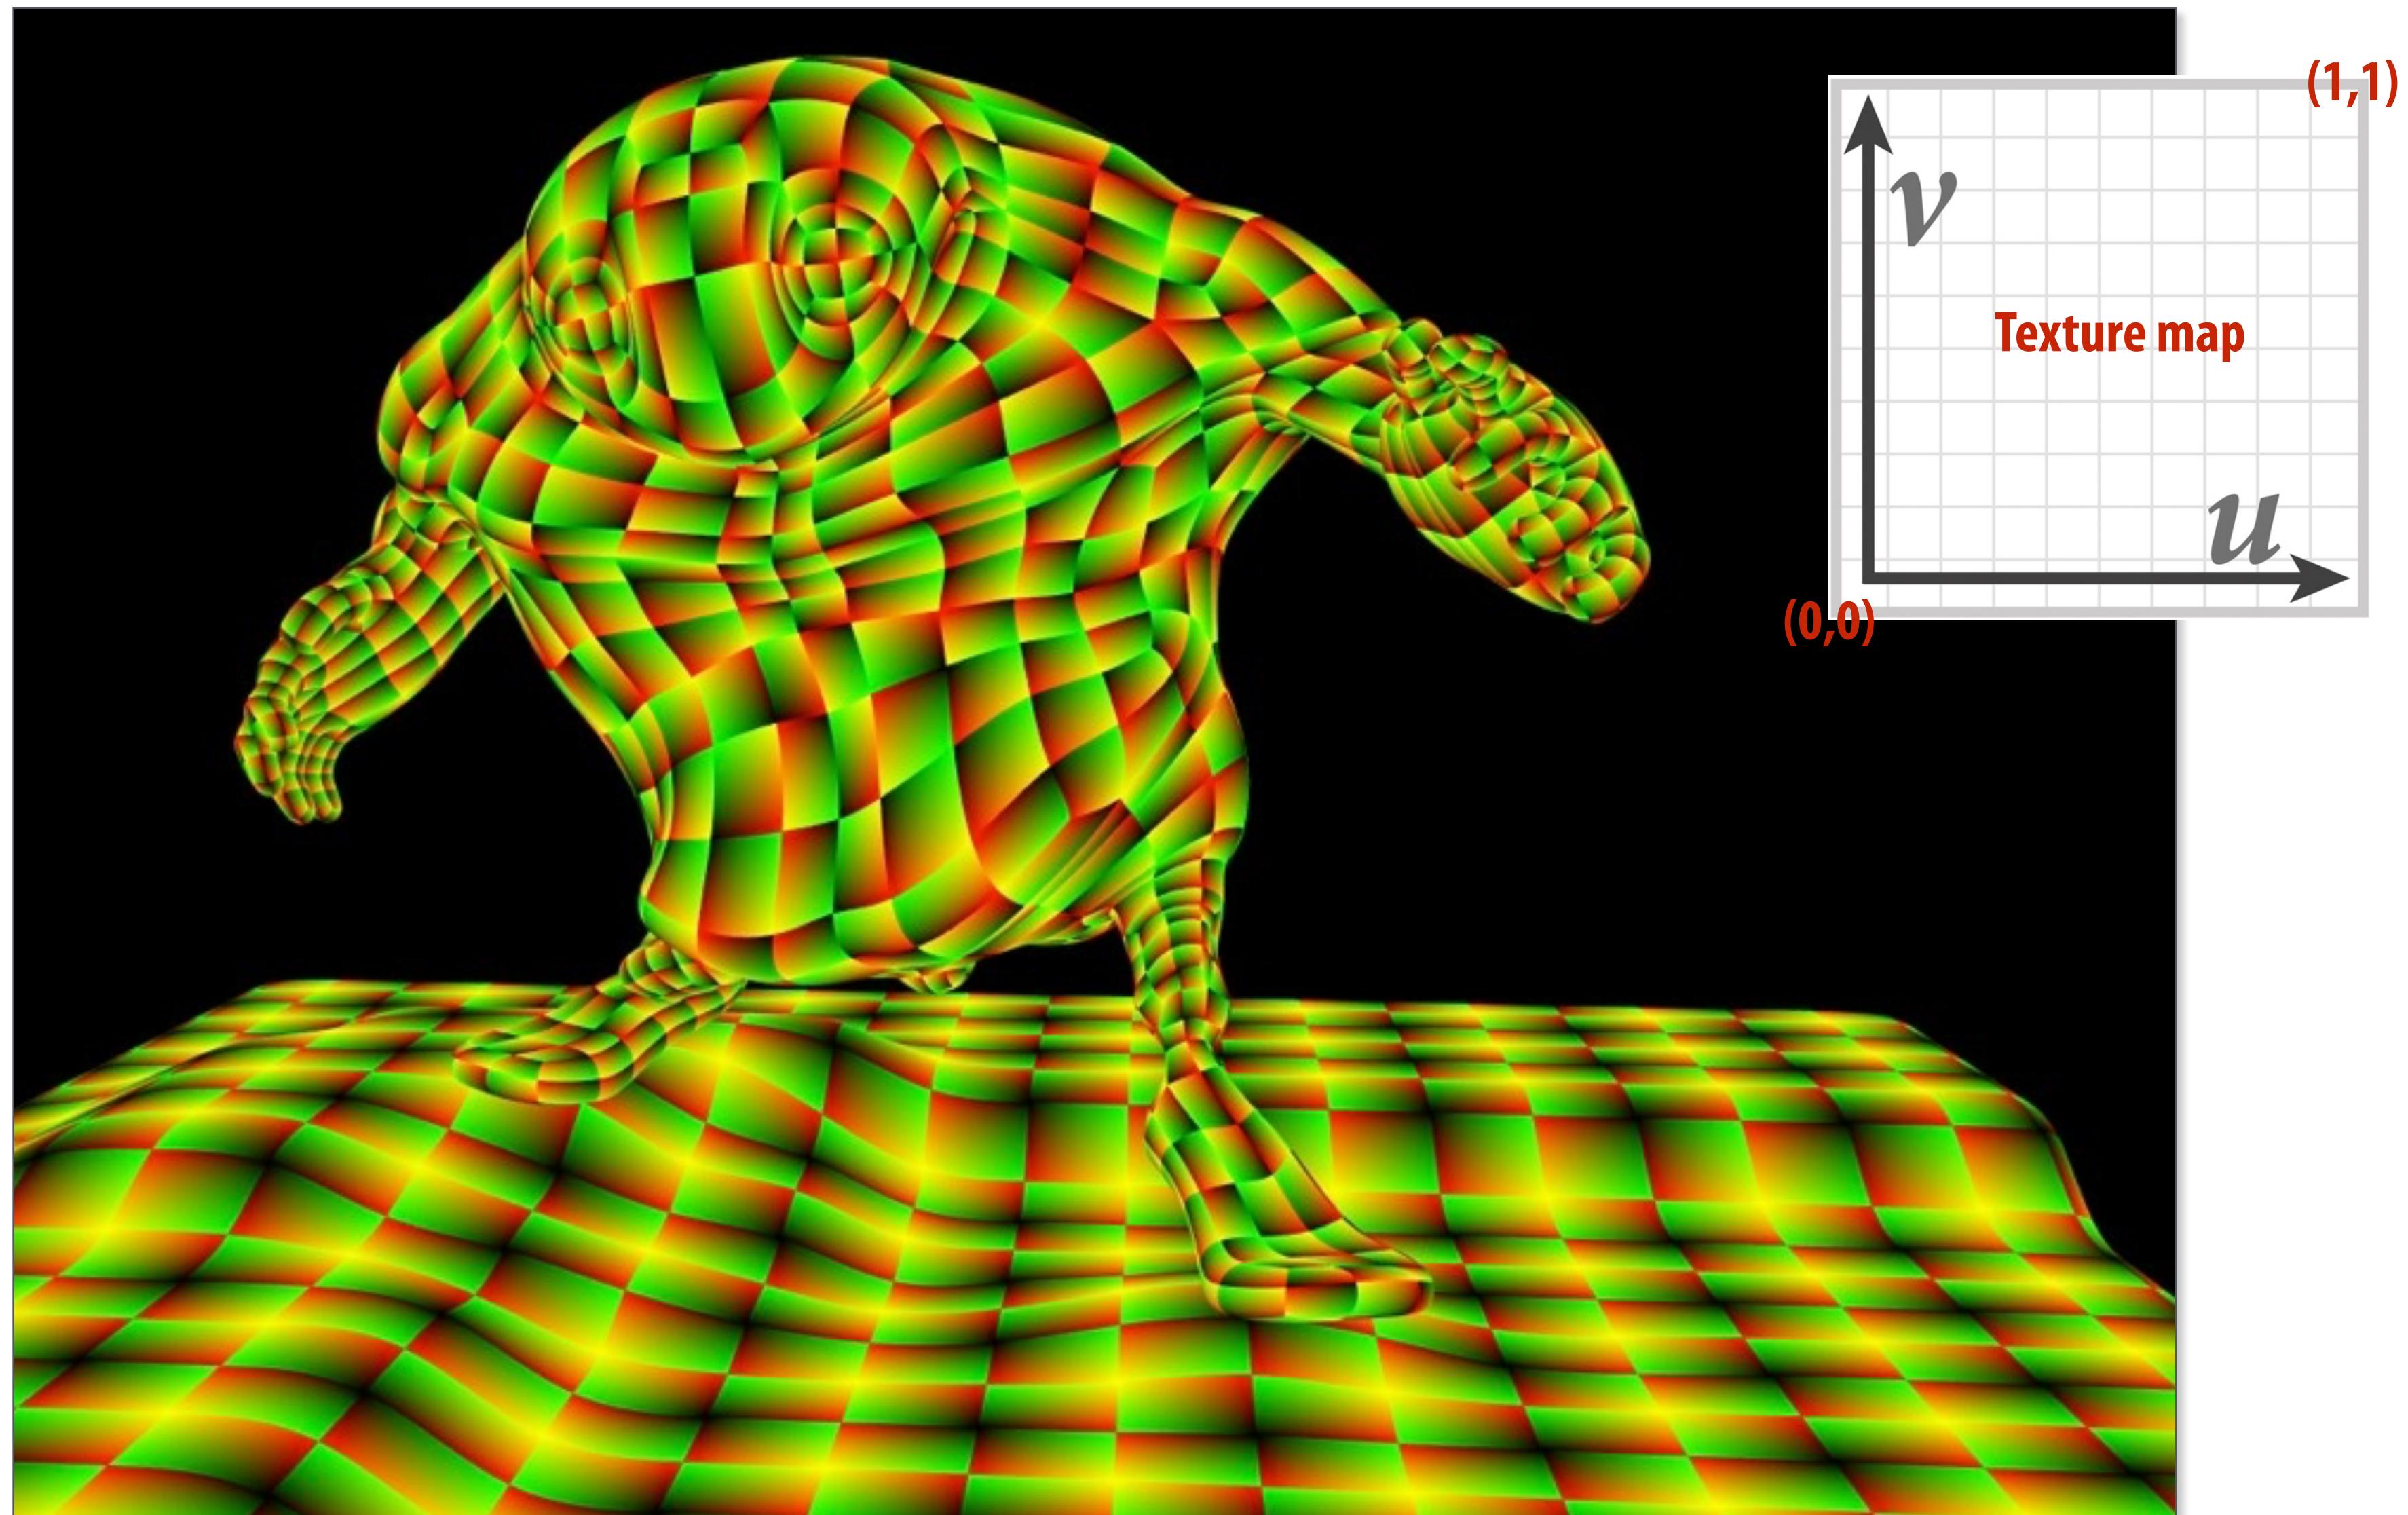
■ Basic algorithm for mapping texture to surface:

- For each color sample location (X, Y)
 - Interpolate U and V coordinates across triangle to get value at (X, Y)
 - Sample (evaluate) texture at location given by (U, V)
 - Set color of surface point to sampled texture value



Texture coordinate visualization

Defines mapping from point on surface to point (uv) in texture domain



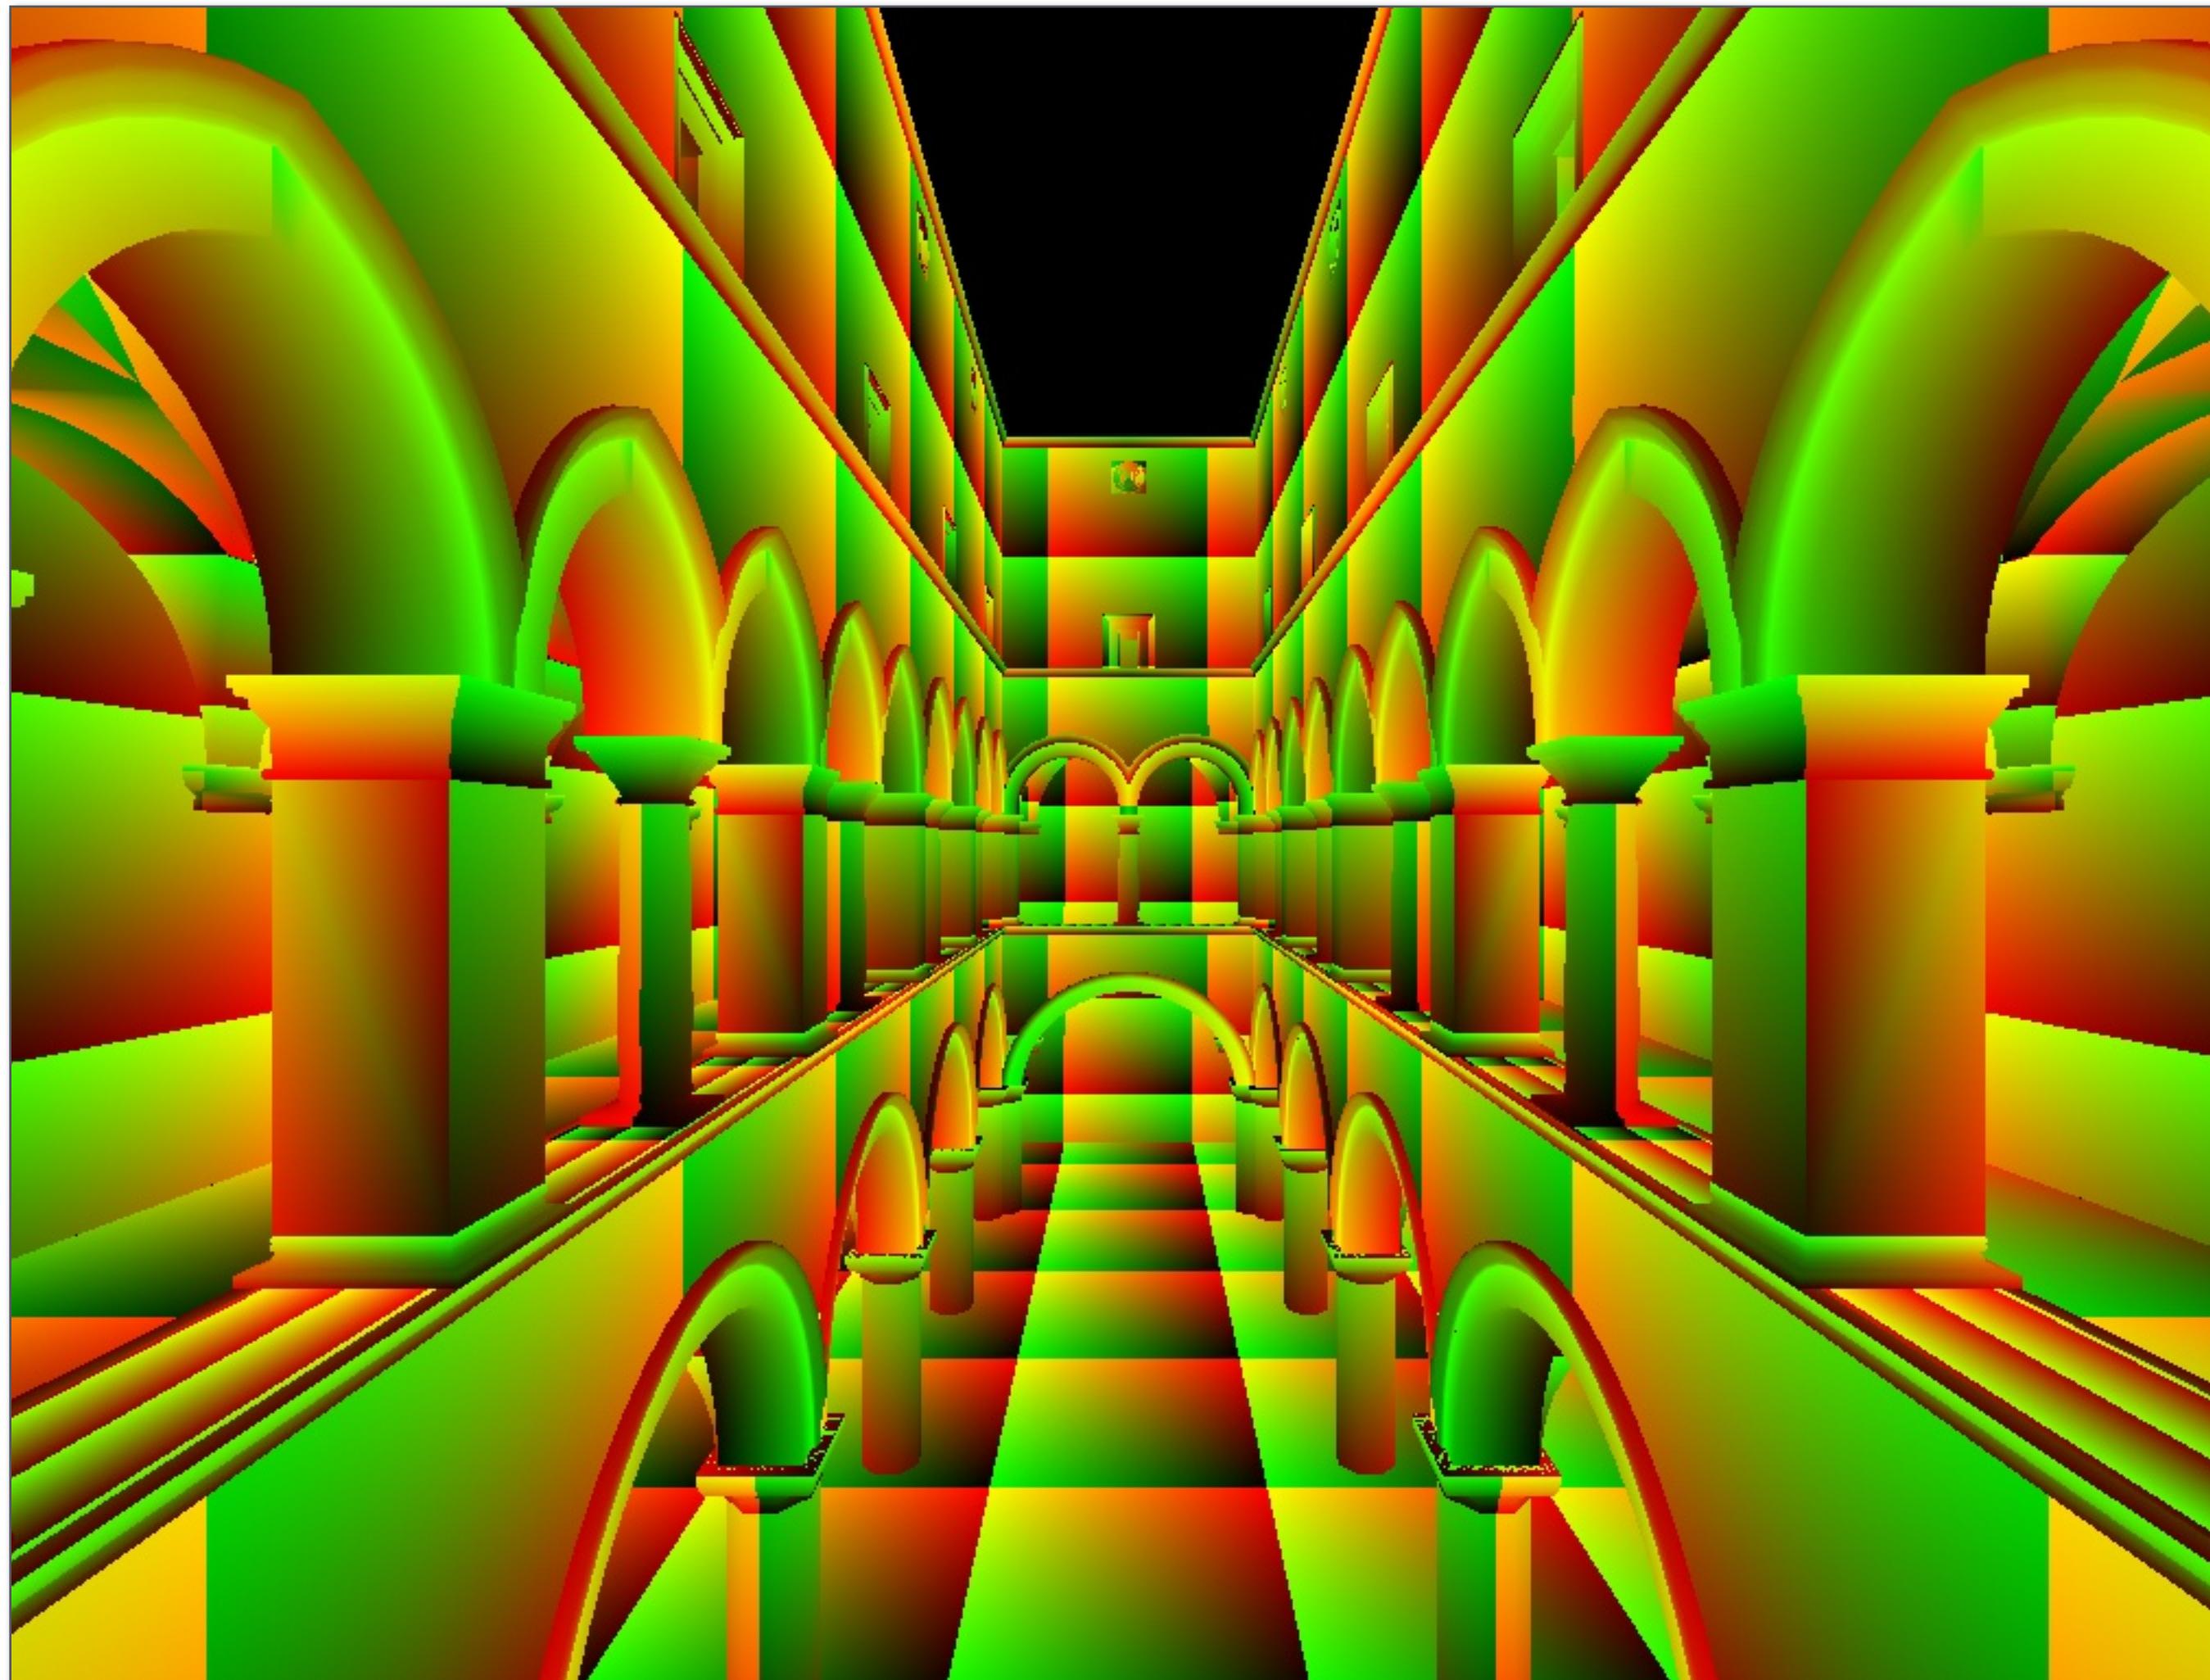
Red channel = u , Green channel = v

So $uv=(0,0)$ is black, $uv=(1,1)$ is yellow

Rendered result

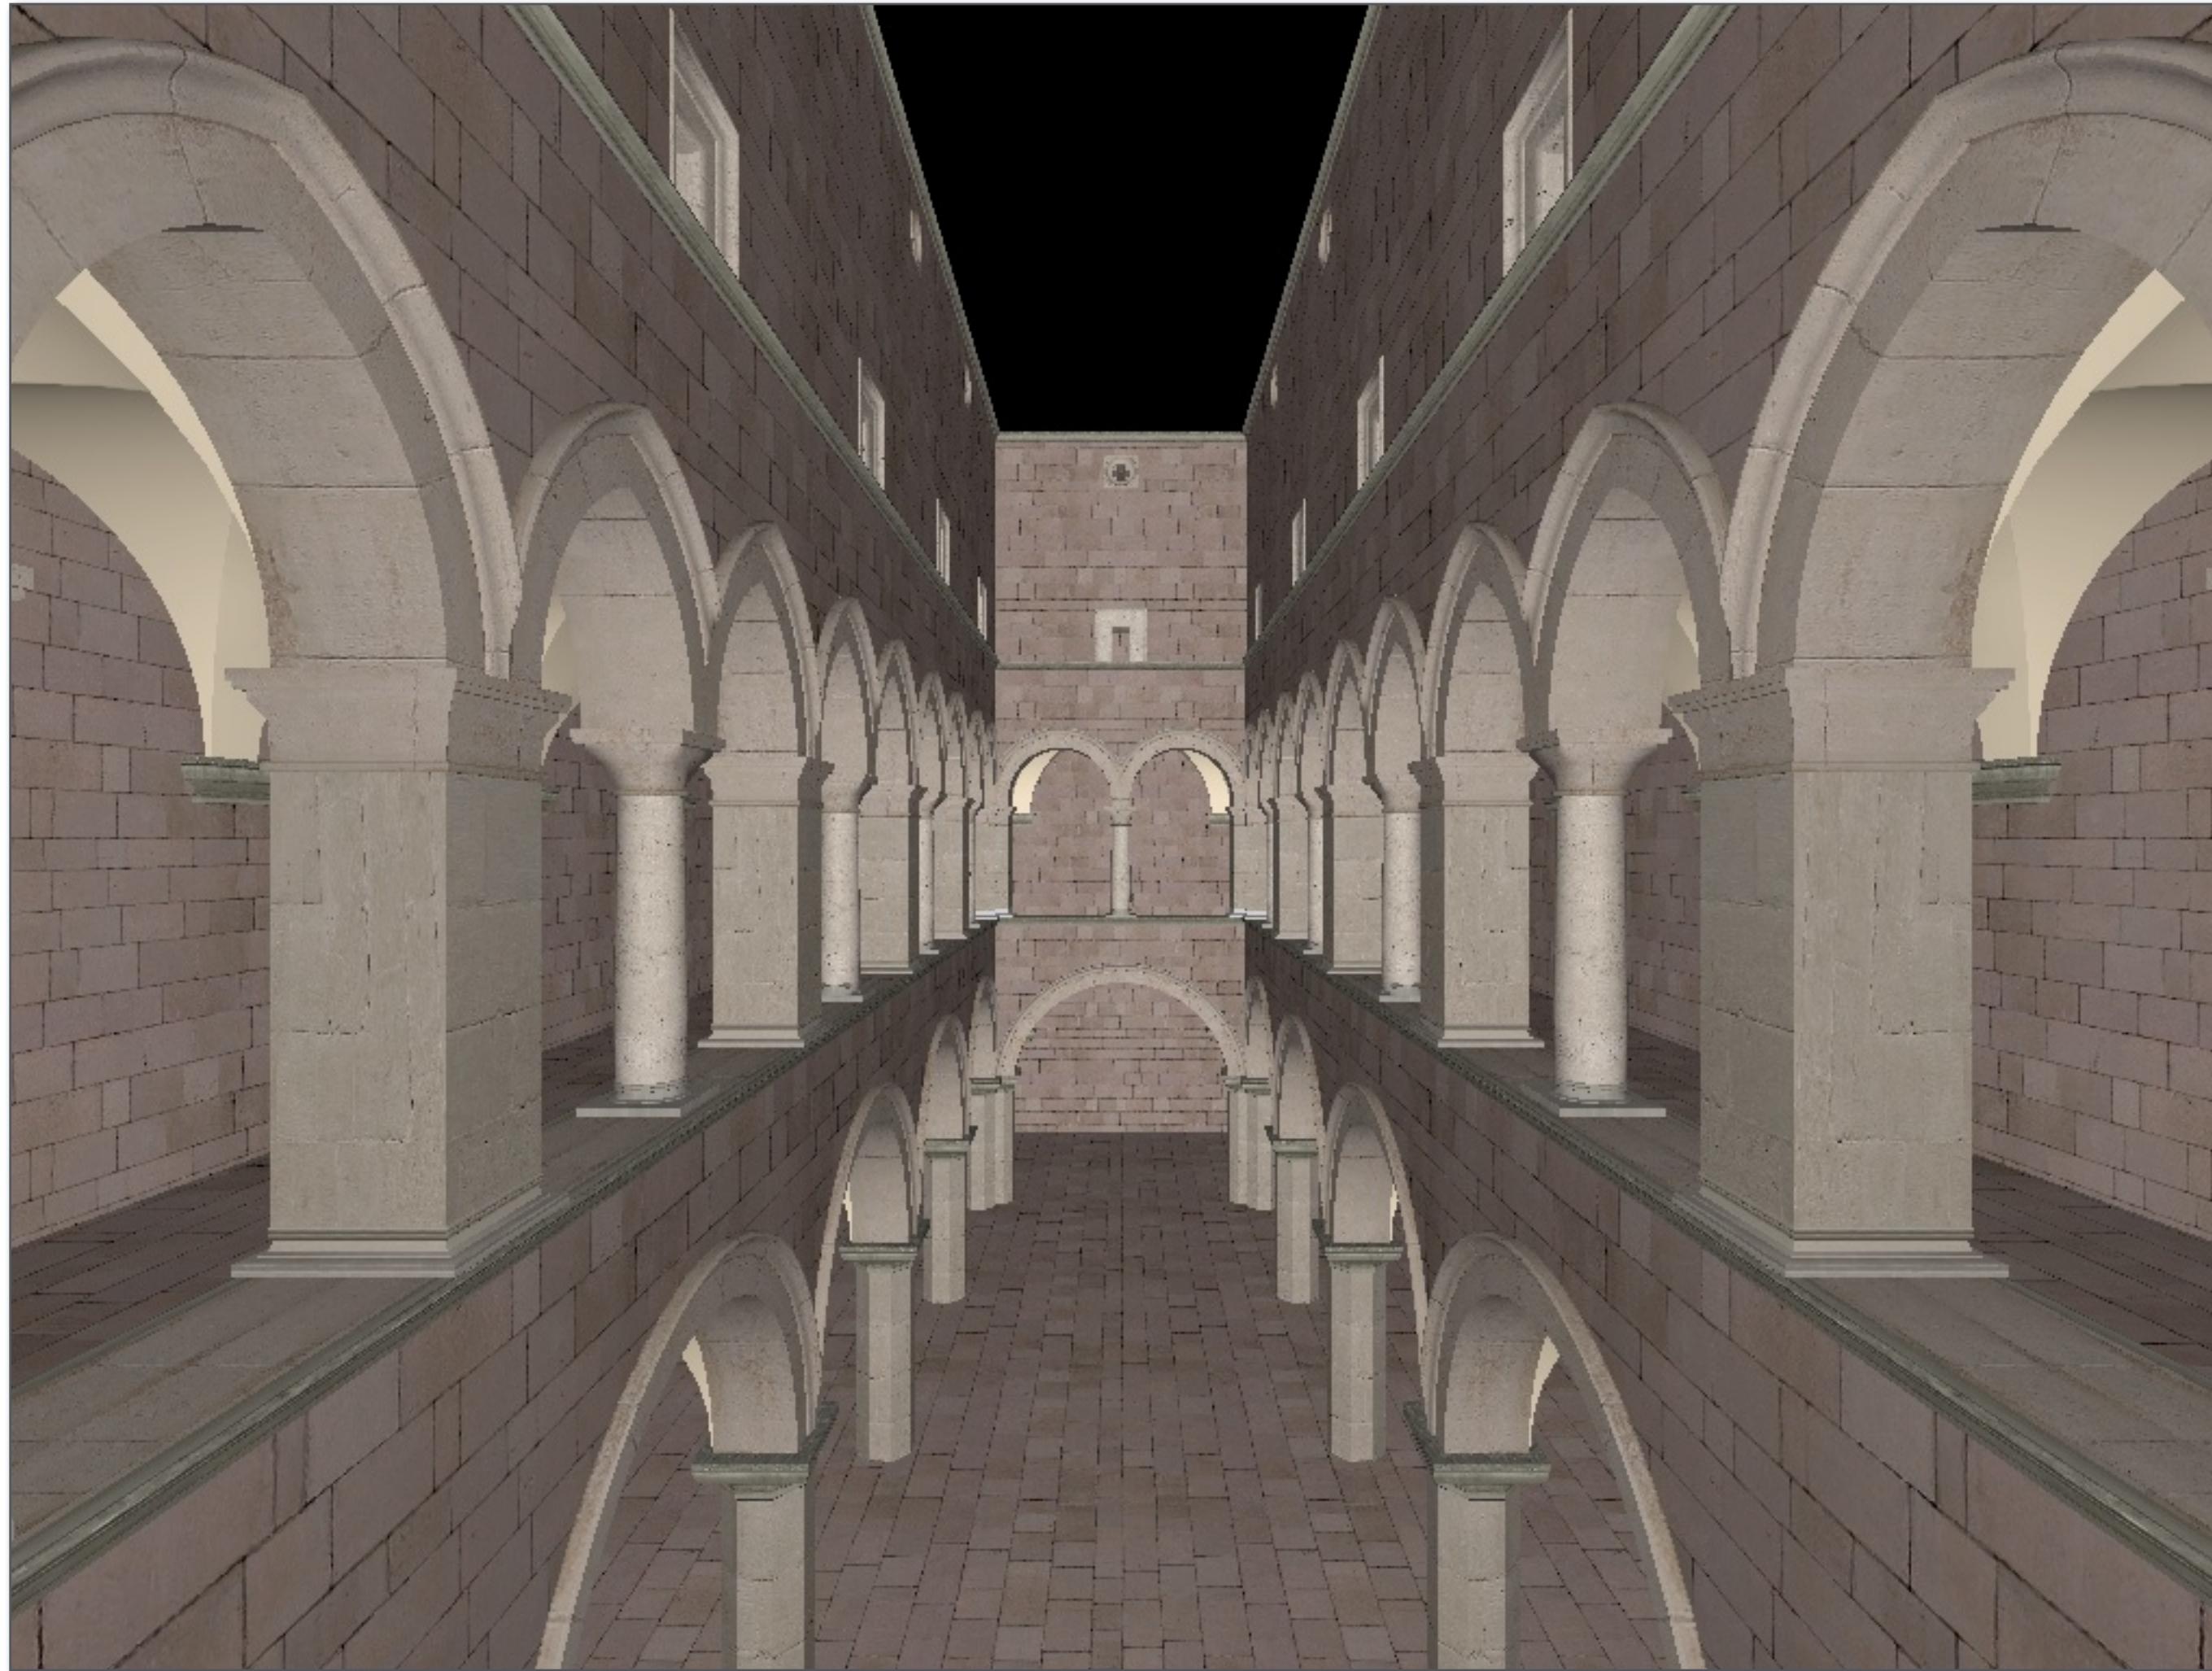


Visualization of texture coordinates

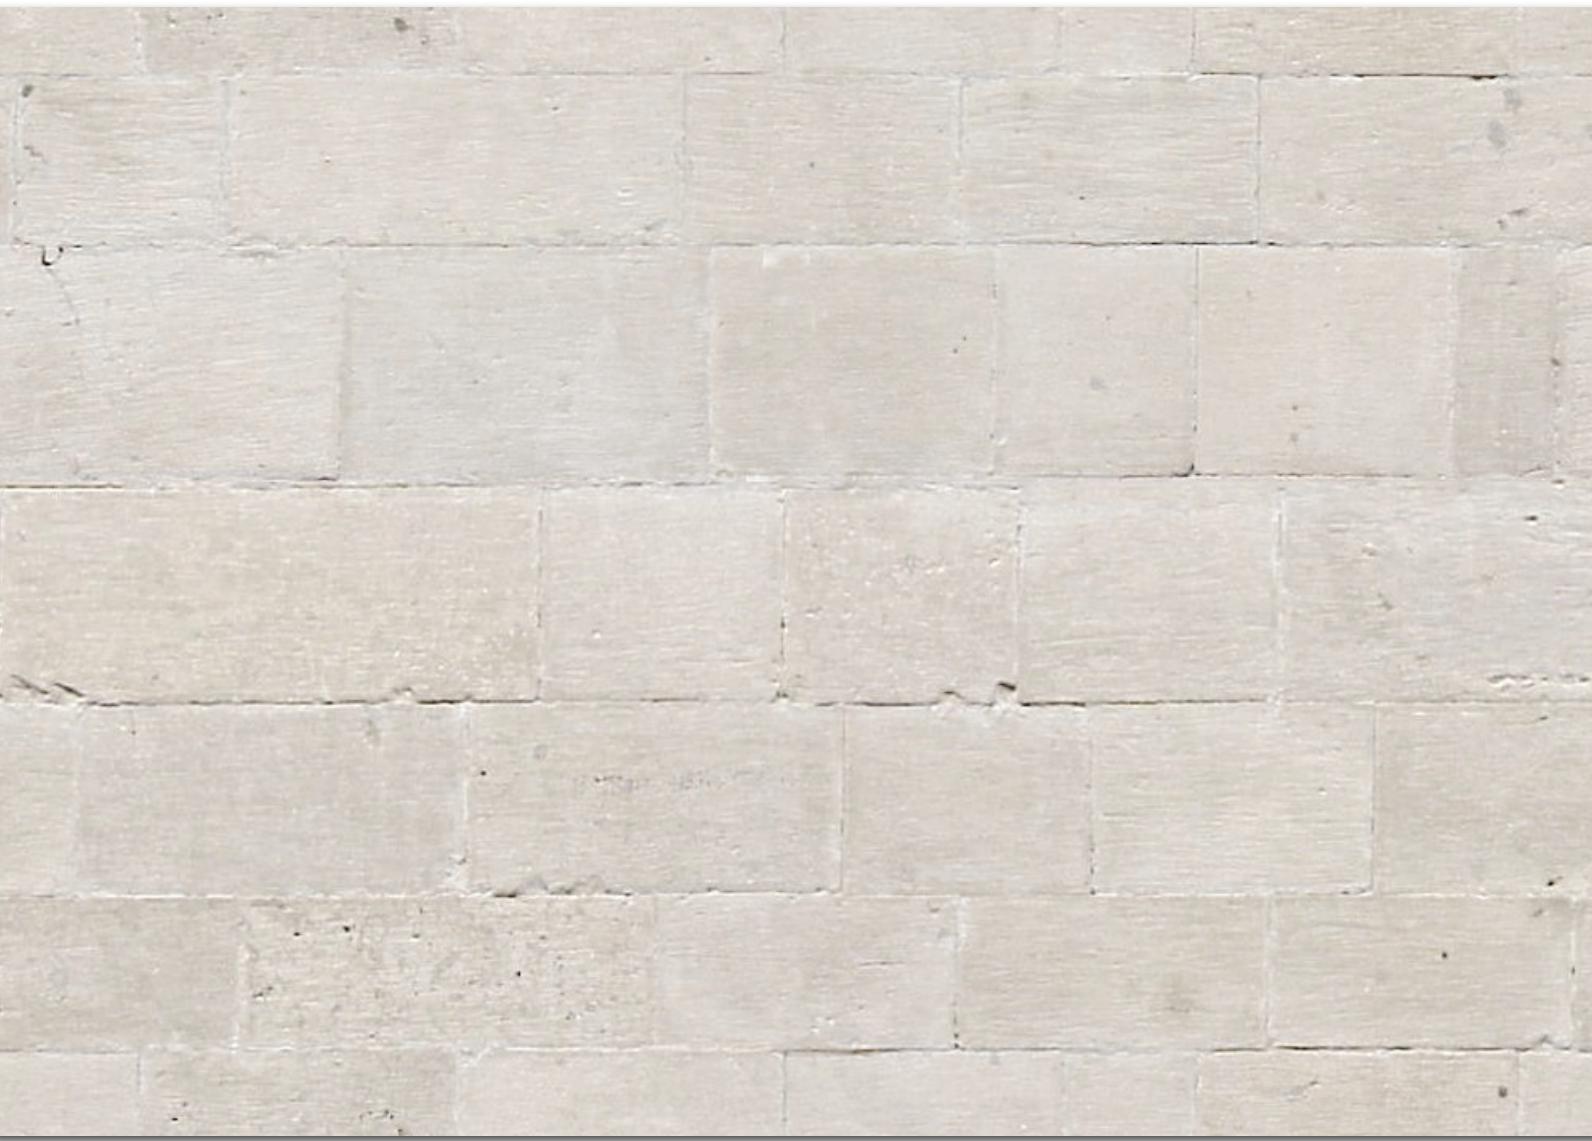
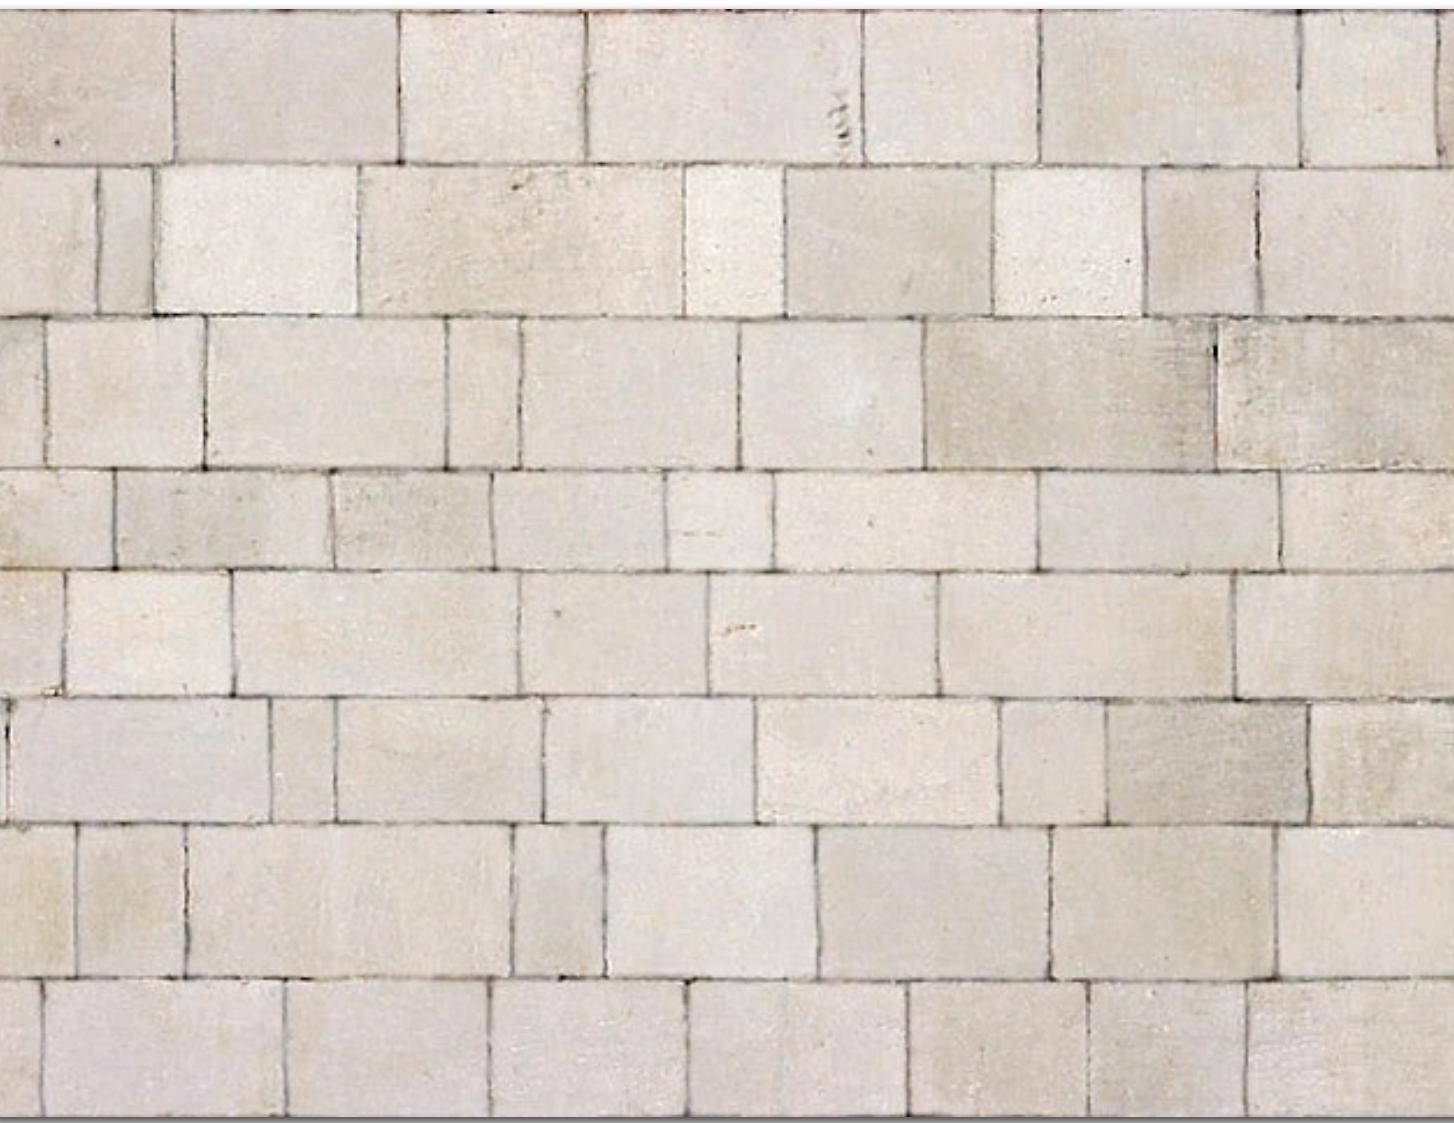


Notice texture coordinates repeat over surface.

Example textured scene

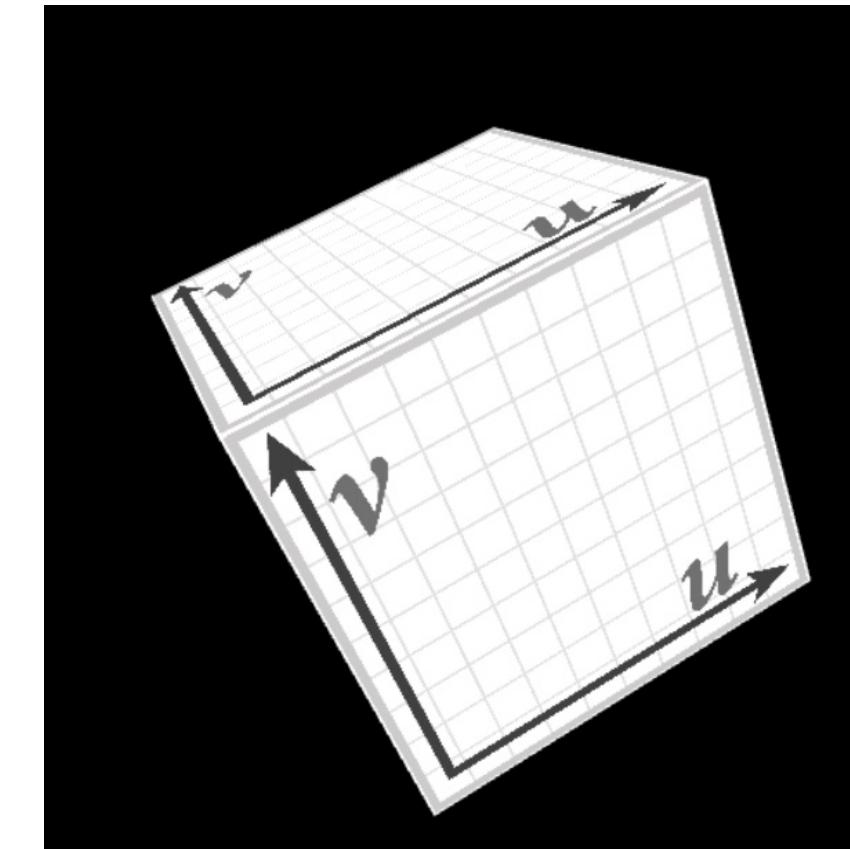
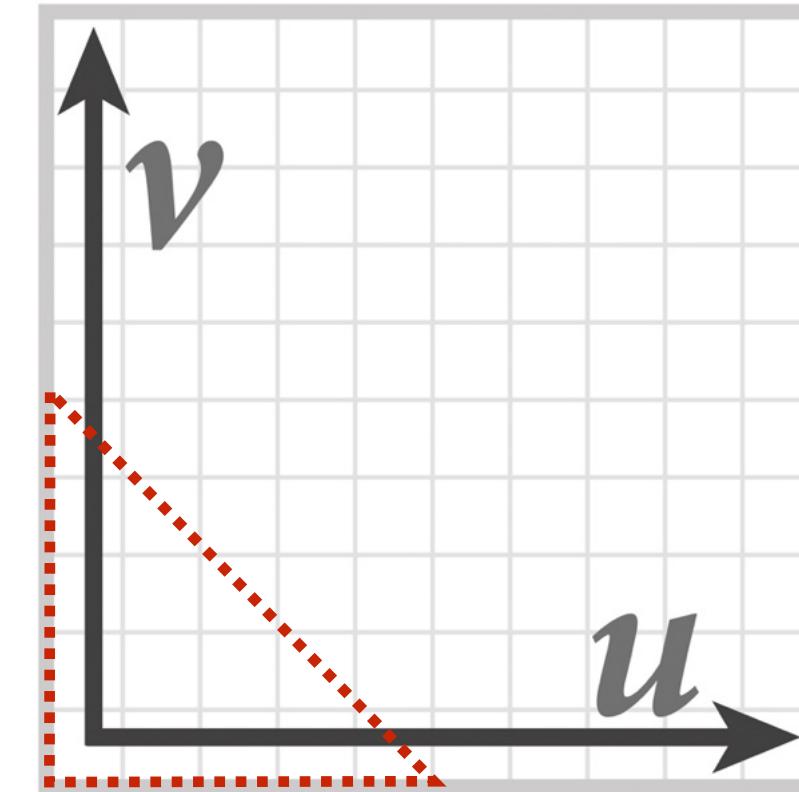
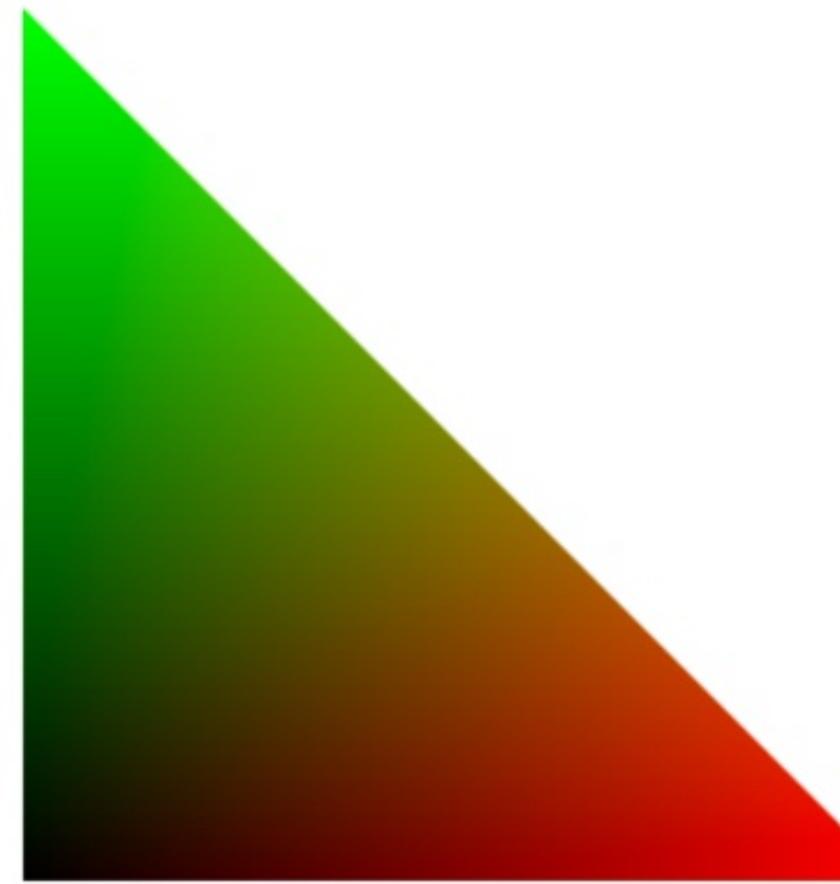


Example textures used in previous scene



Texture mapping: basic algorithm (ray

- **Basic algorithm for mapping texture image onto a surface:**
 - For each camera ray, find surface intersection point P with triangle
 - Interpolate U and V texture coordinates across triangle to get texture coordinate value at P
 - Sample texture map at location (U, V)
 - Set output image sample color to sampled texture value

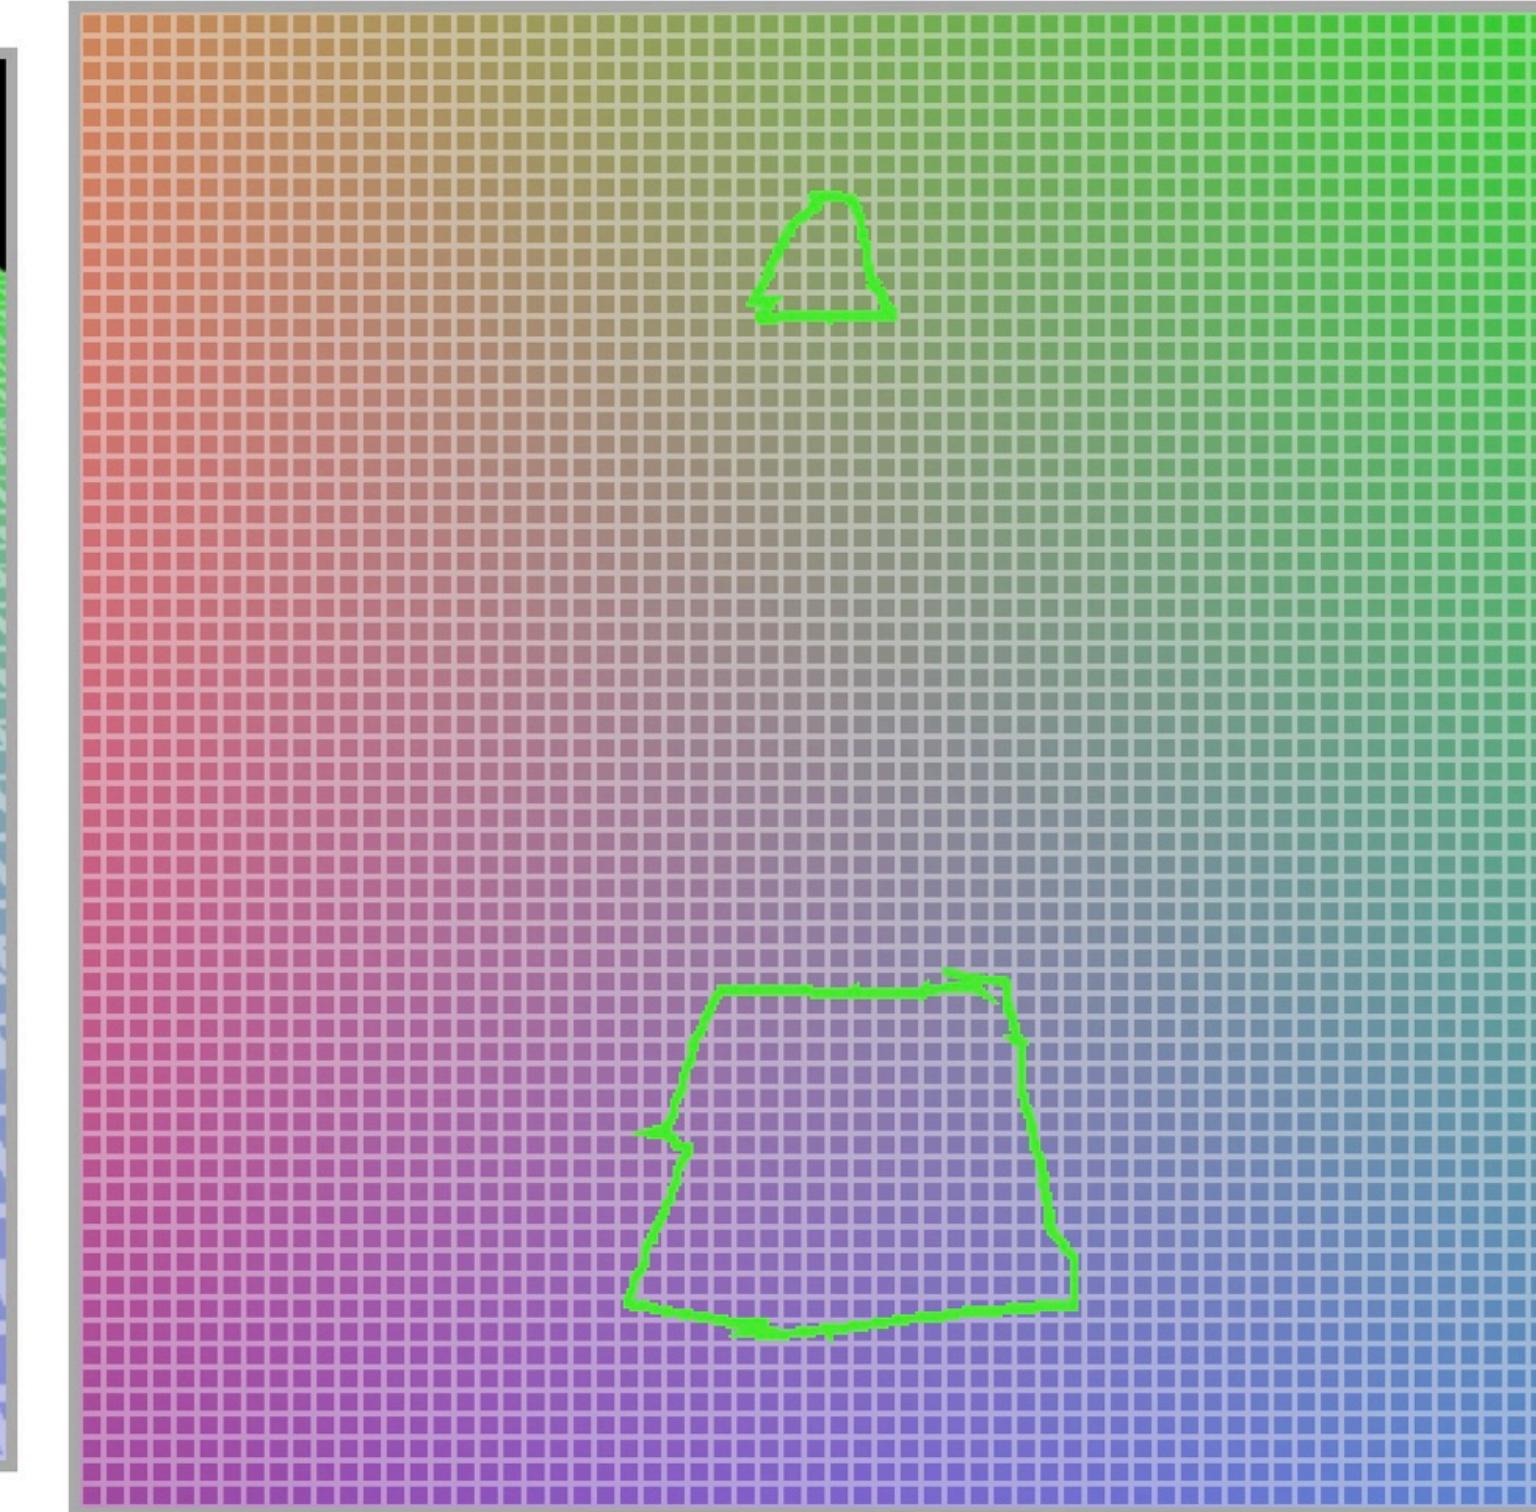
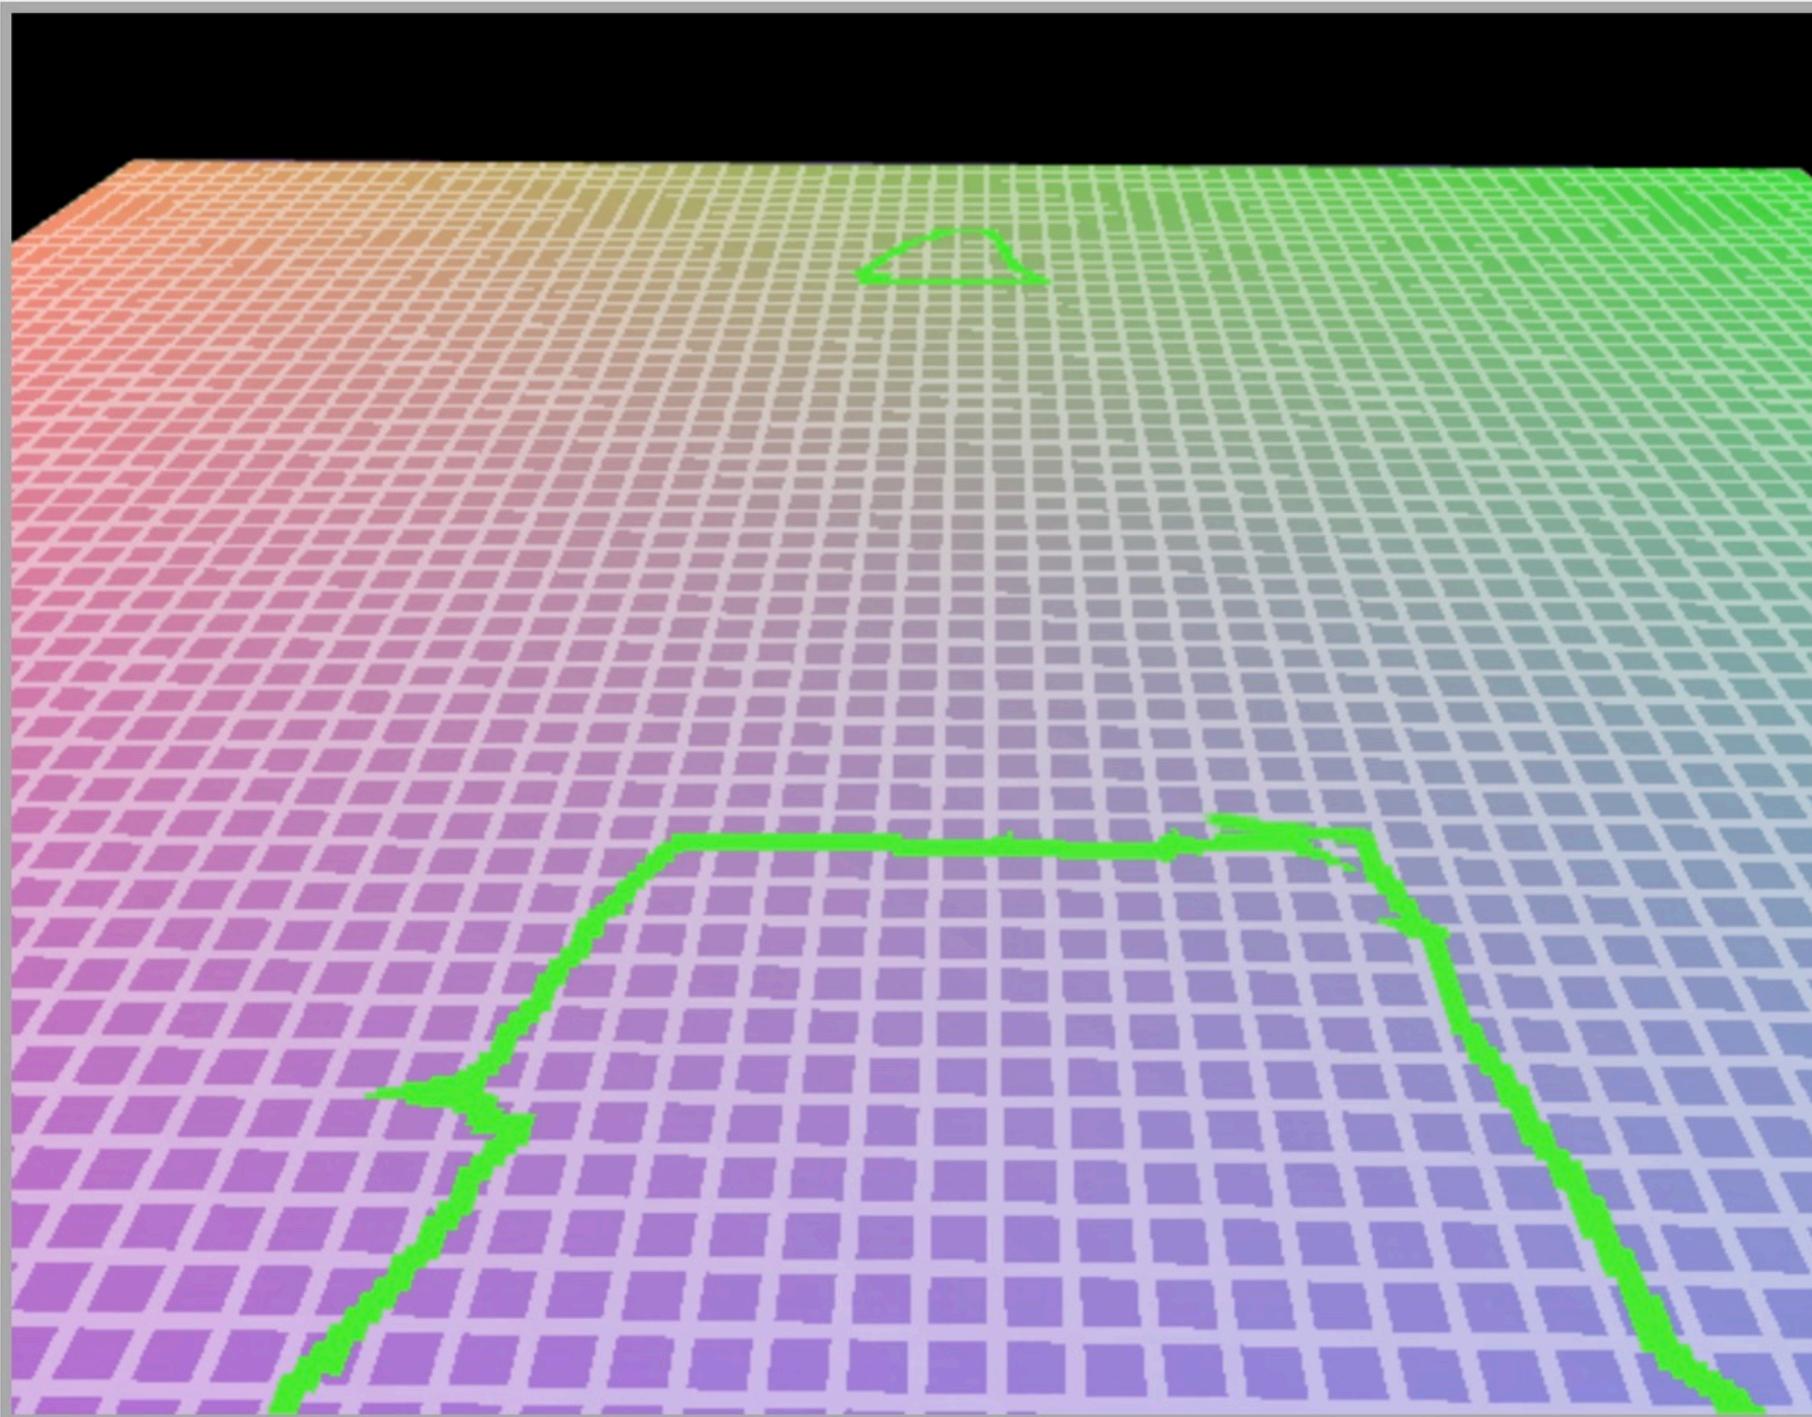


Demo (by Katie Detkar)

<https://katie.su.domains/webgl/index.html>

Image warp through texturing and projection

Below are two images. The first is a 3D rendering of a textured model, and the second is a 2D visualization of texture space. You can compare the first compared to the second during the transformations that take it from object space to screen space.



Object

Model type

Rotation around horizontal

Rotation around vertical

Object scale

Base texture

Base

Texture mapping

Nearest neighbor

Sonification

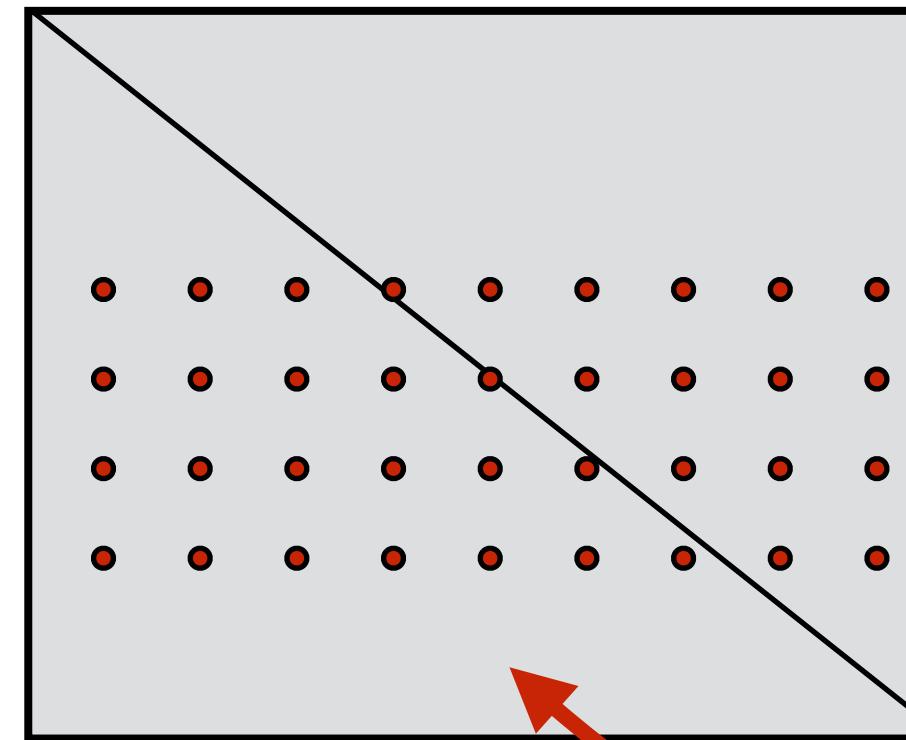
Listen to the Sonification base texture

Thought experiment

Imagine rendering a texture-mapped quadrilateral onto a 1000x1000 pixel output image

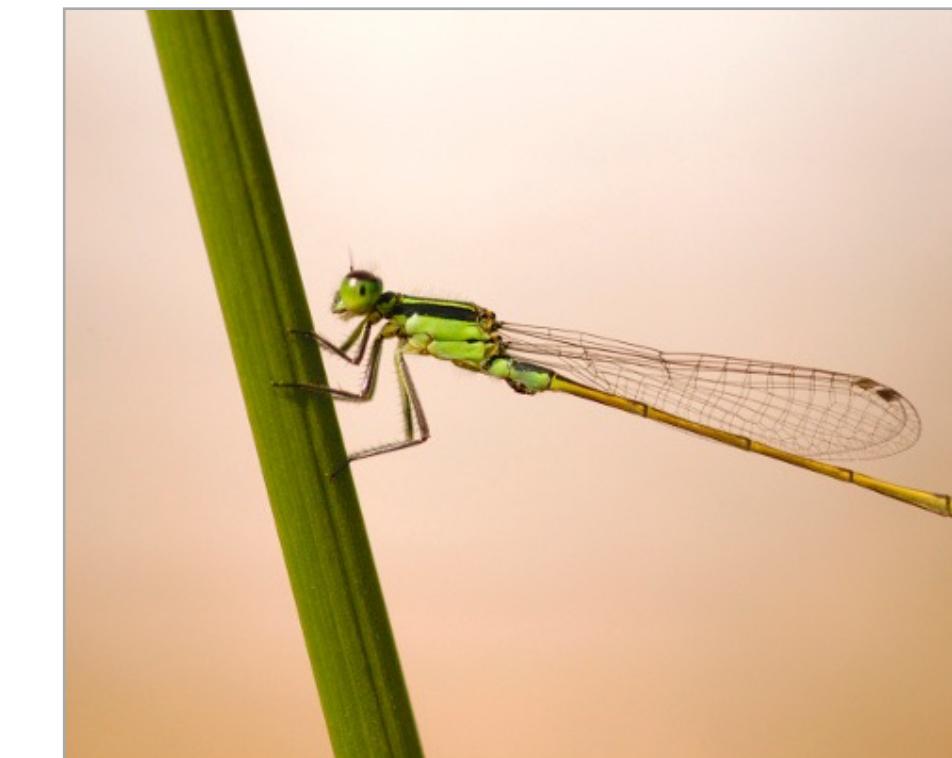
1000 pixels

These red dots are your screen sample locations.

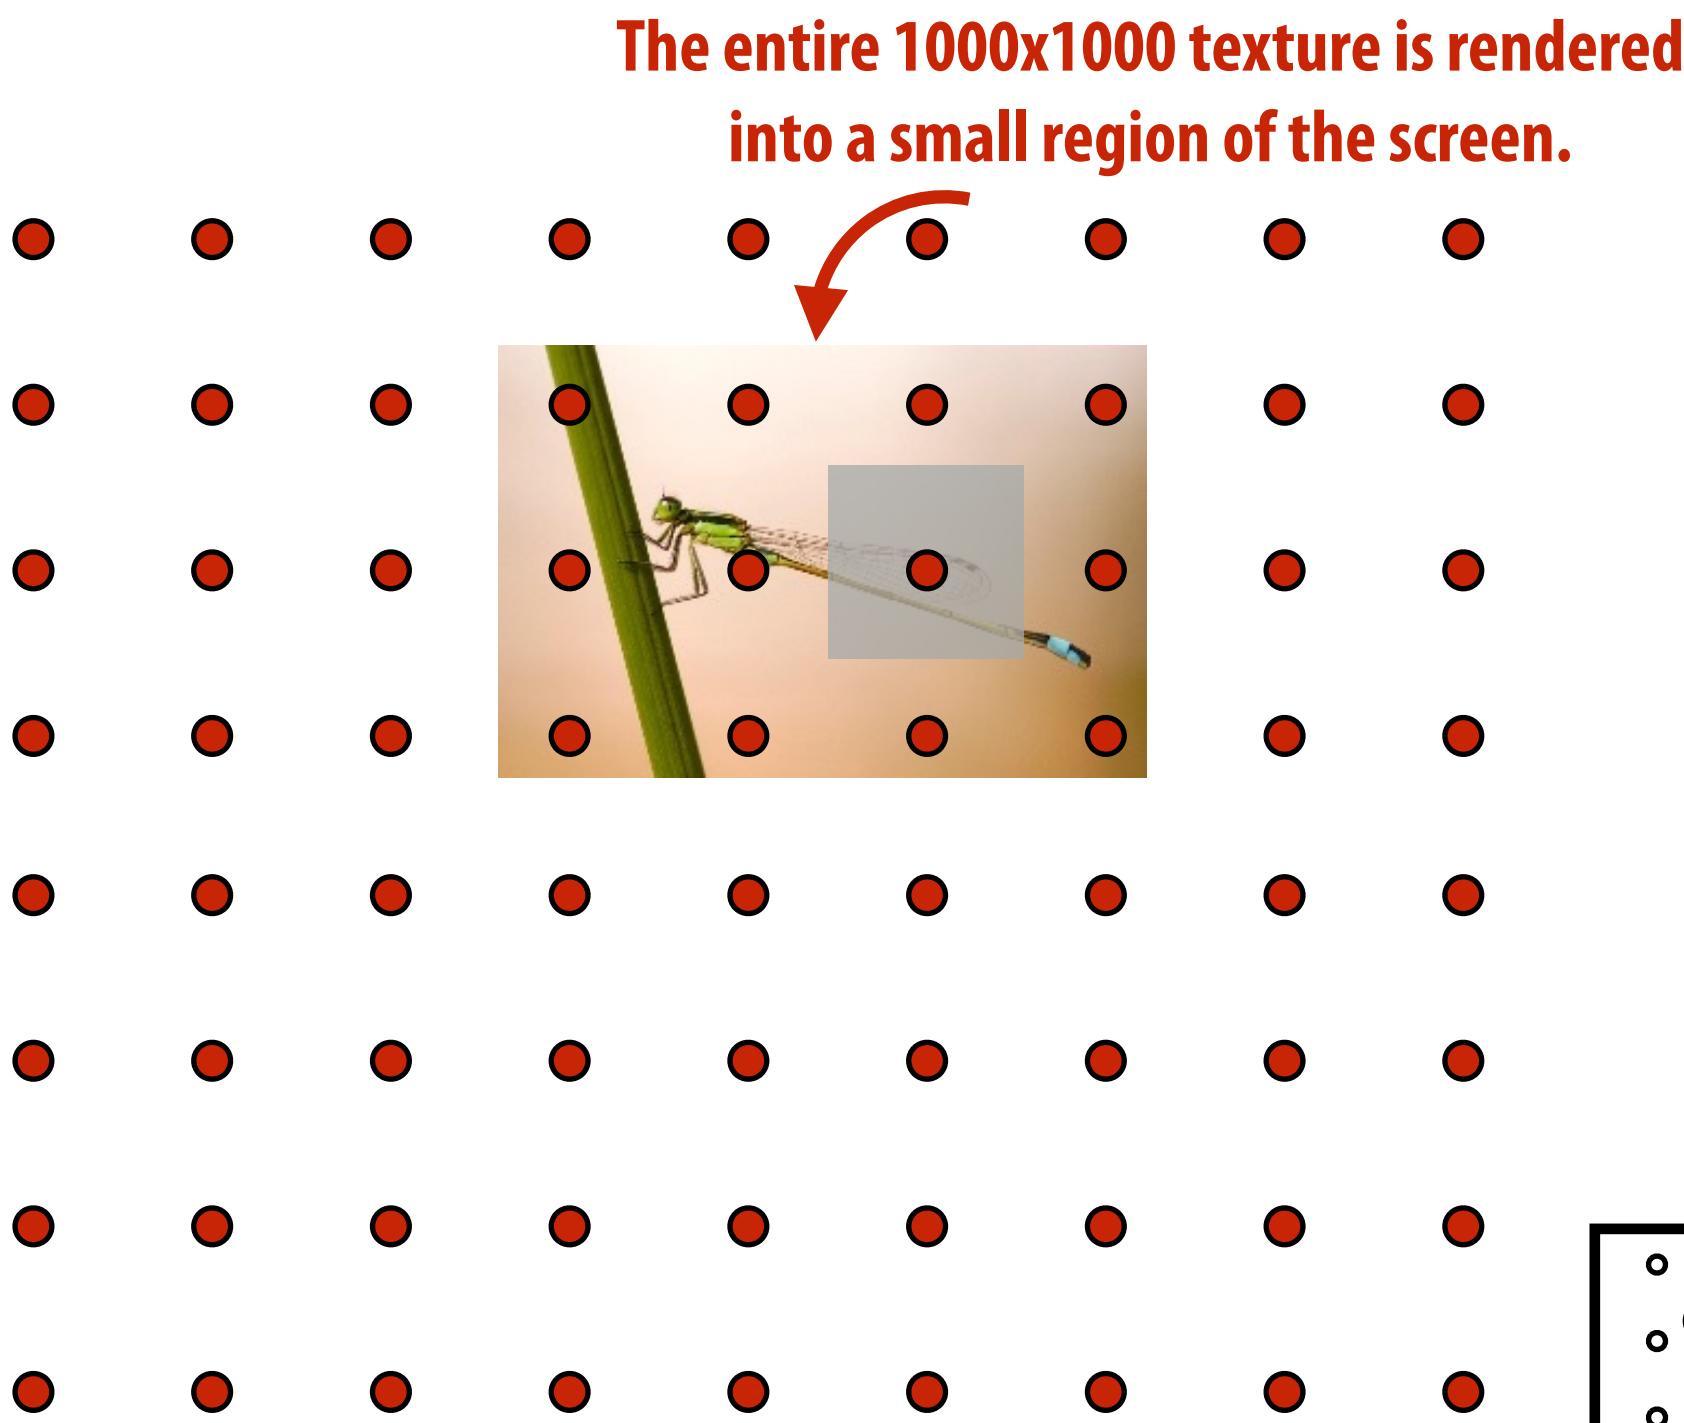


1000 pixels

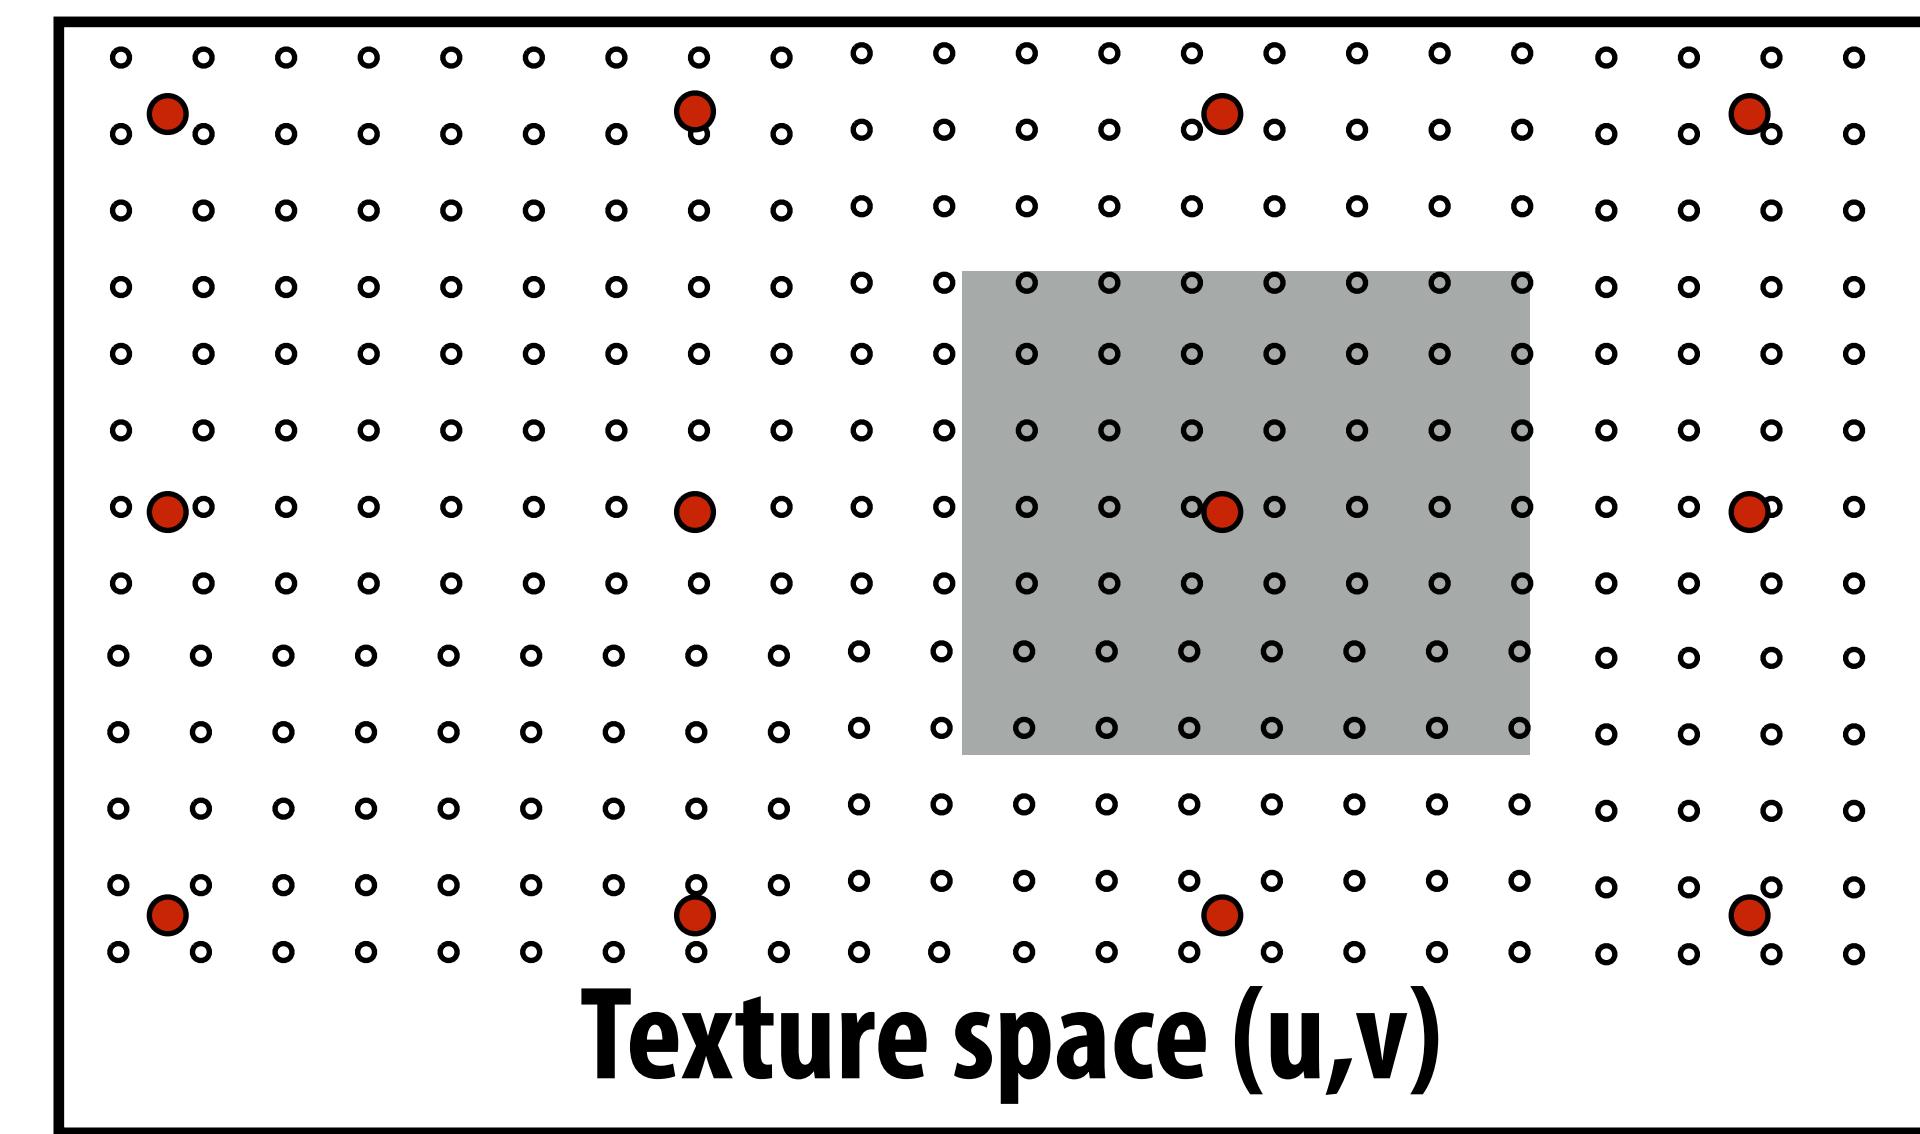
Let's also say the texture image is 1000x1000 as well.



Sampling rate on screen vs in texture: object zoomed out



Texture is “minified” on screen



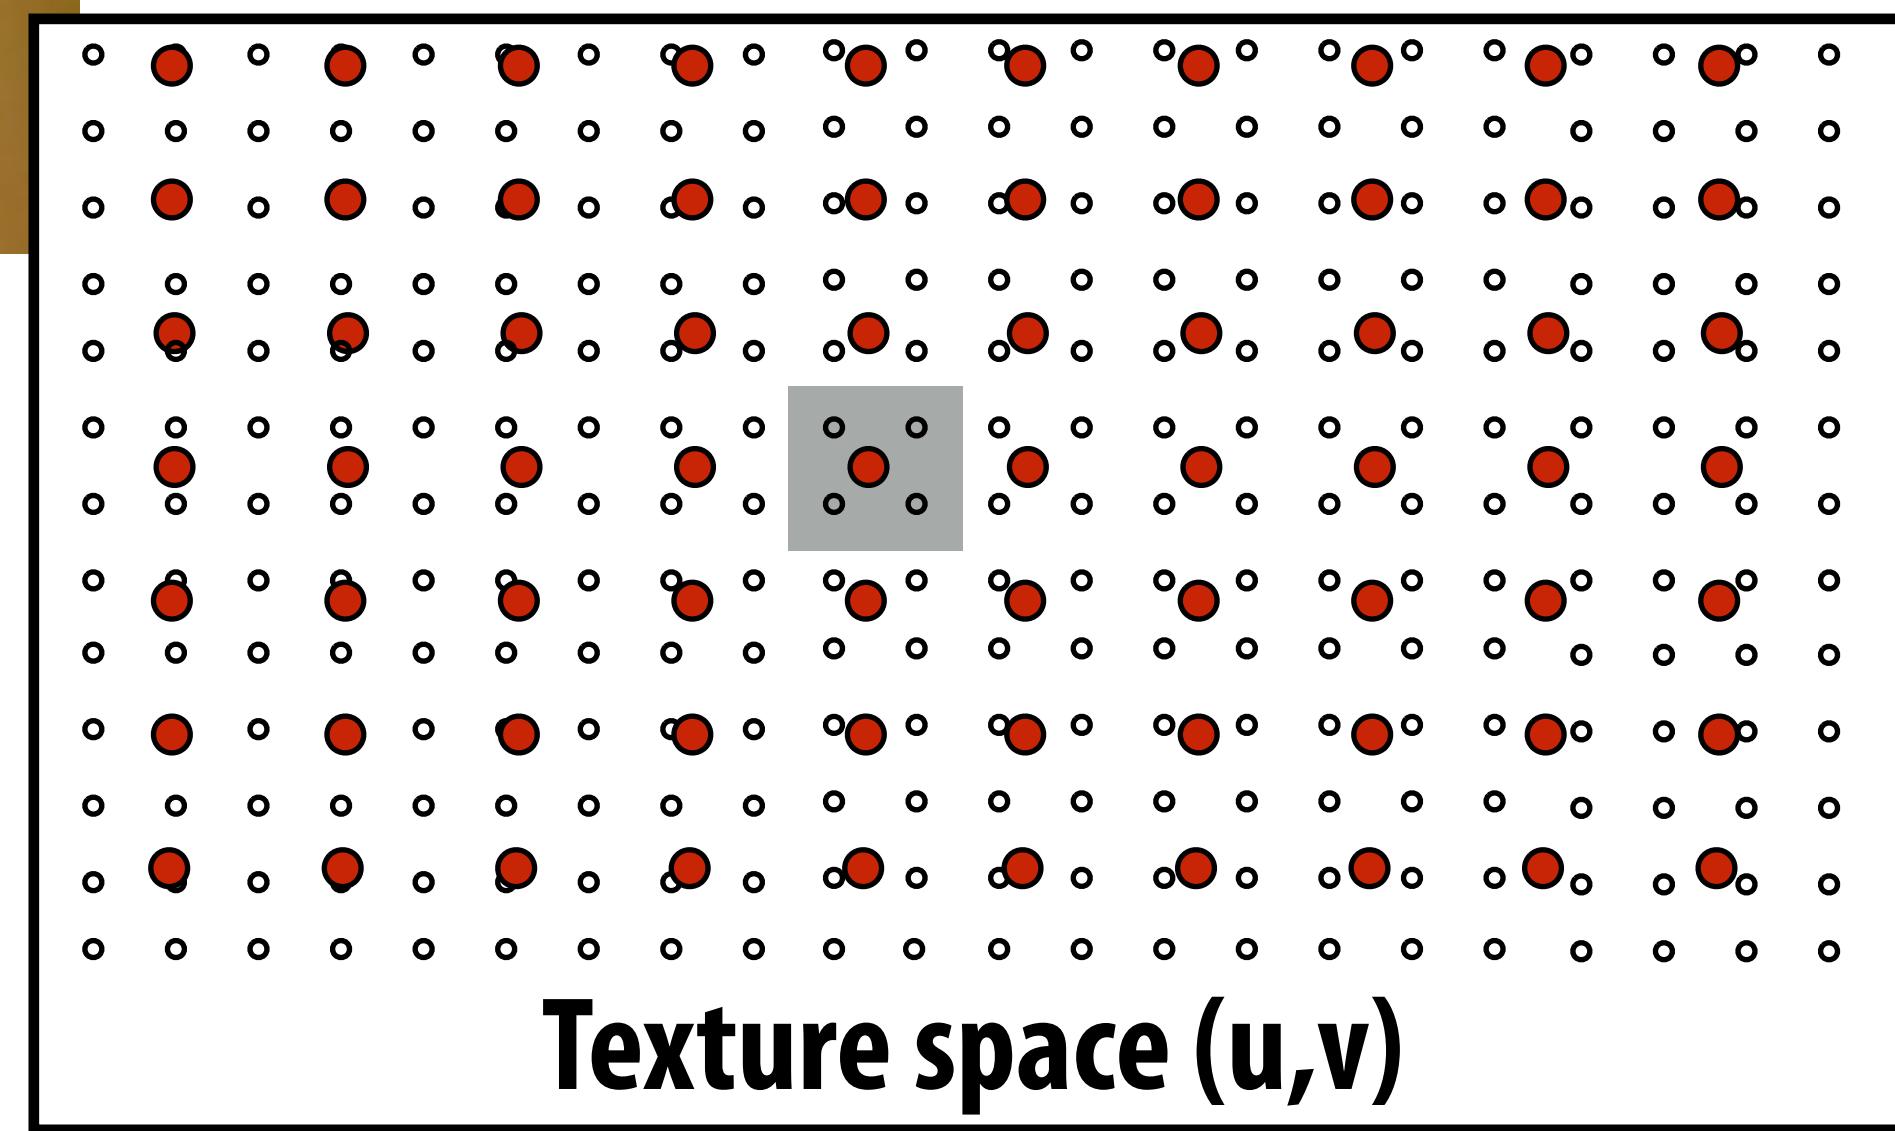
Zooming in...



Red dots = sample positions on screen

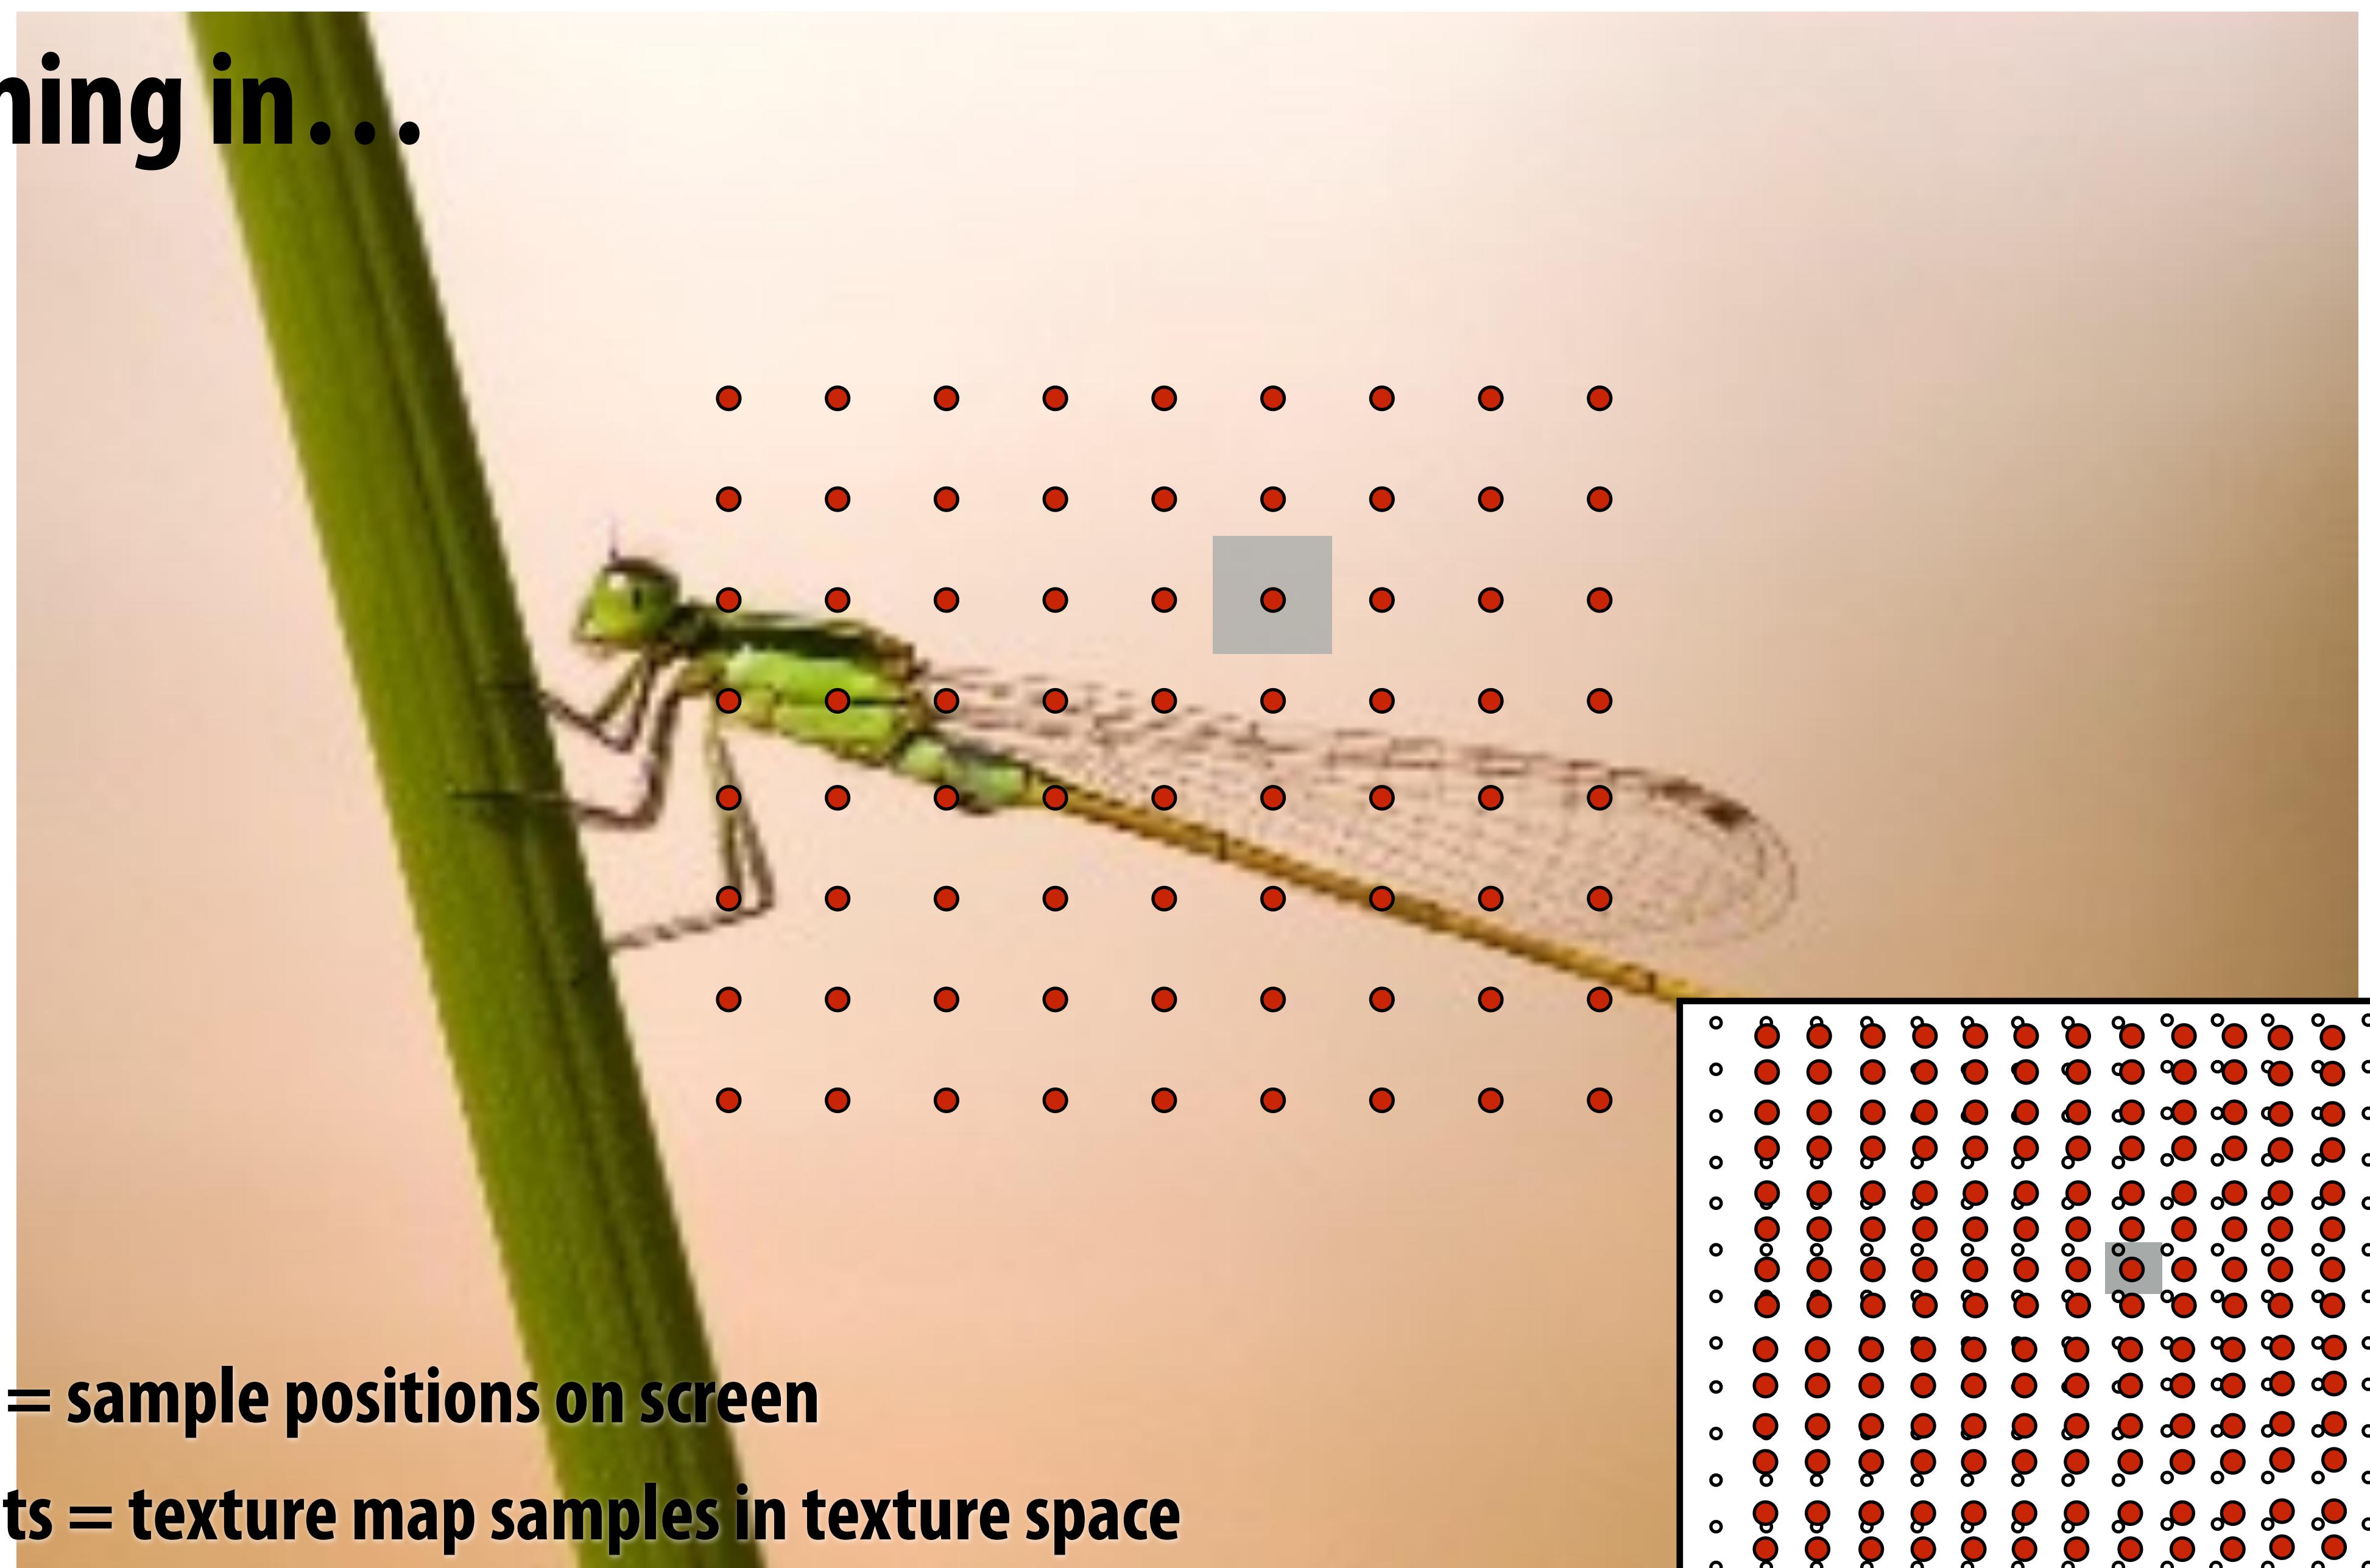
White dots = texture map samples in texture space

Gray square = area of a screen pixel



Texture space (u, v)

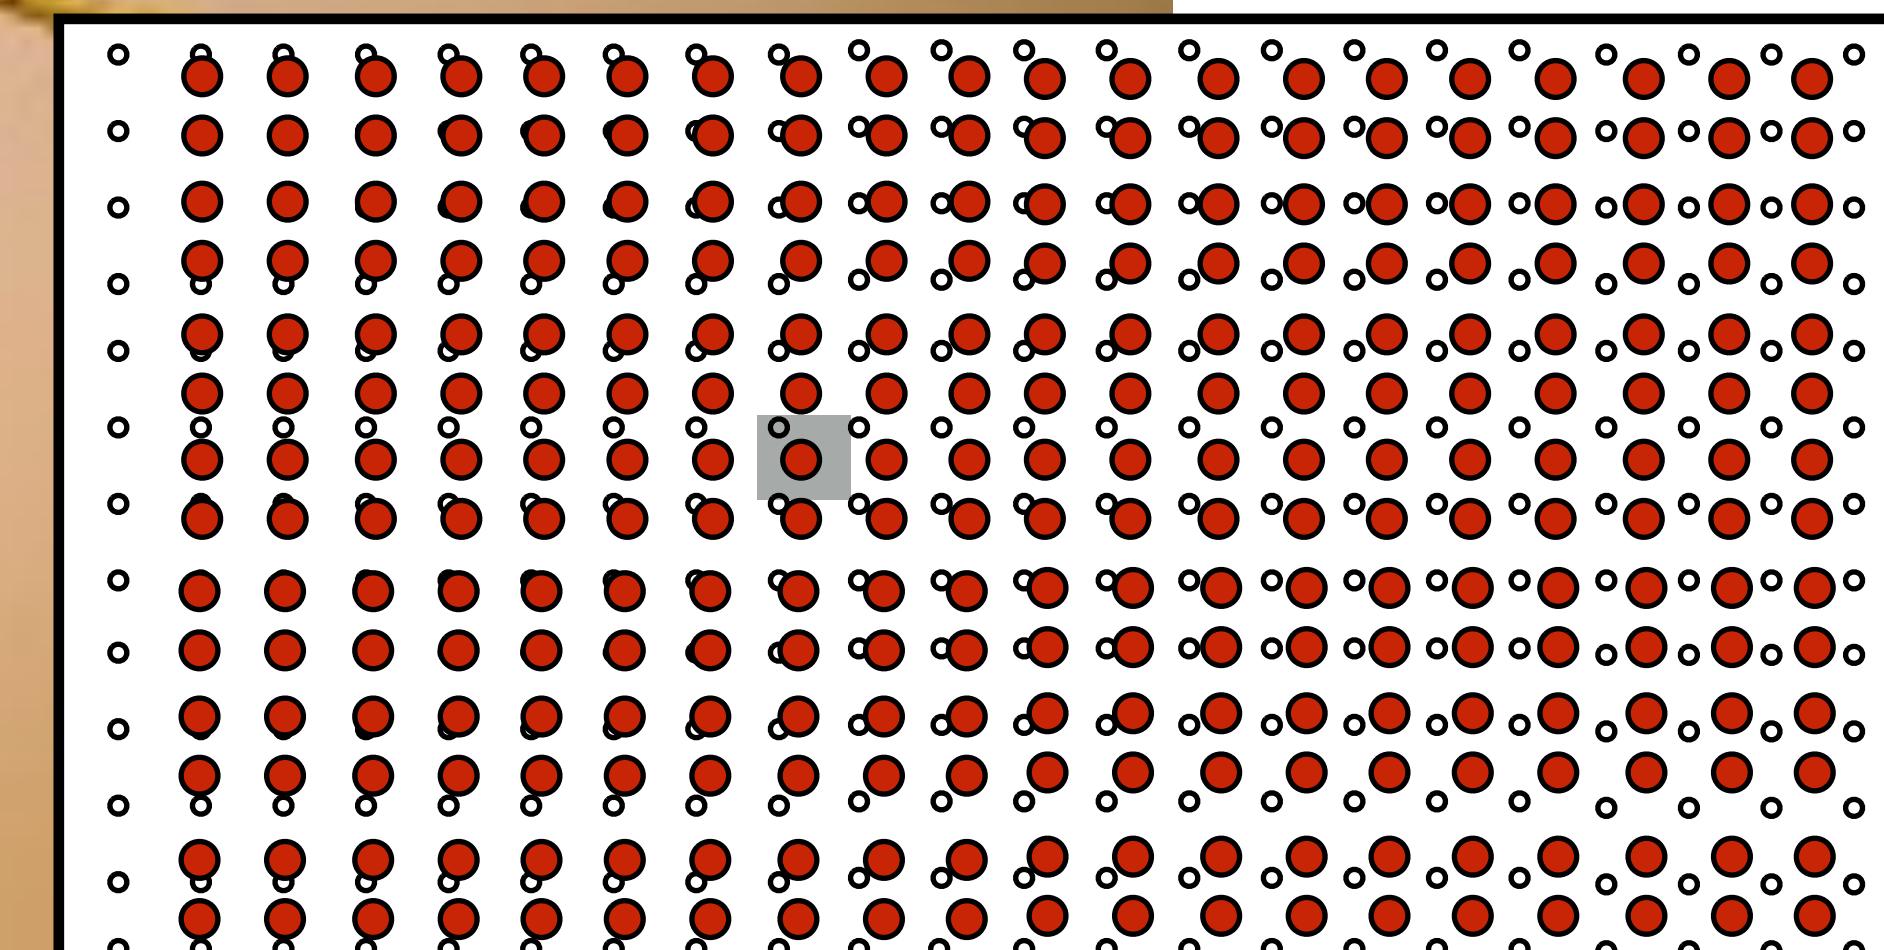
Zooming in...



Red dots = sample positions on screen

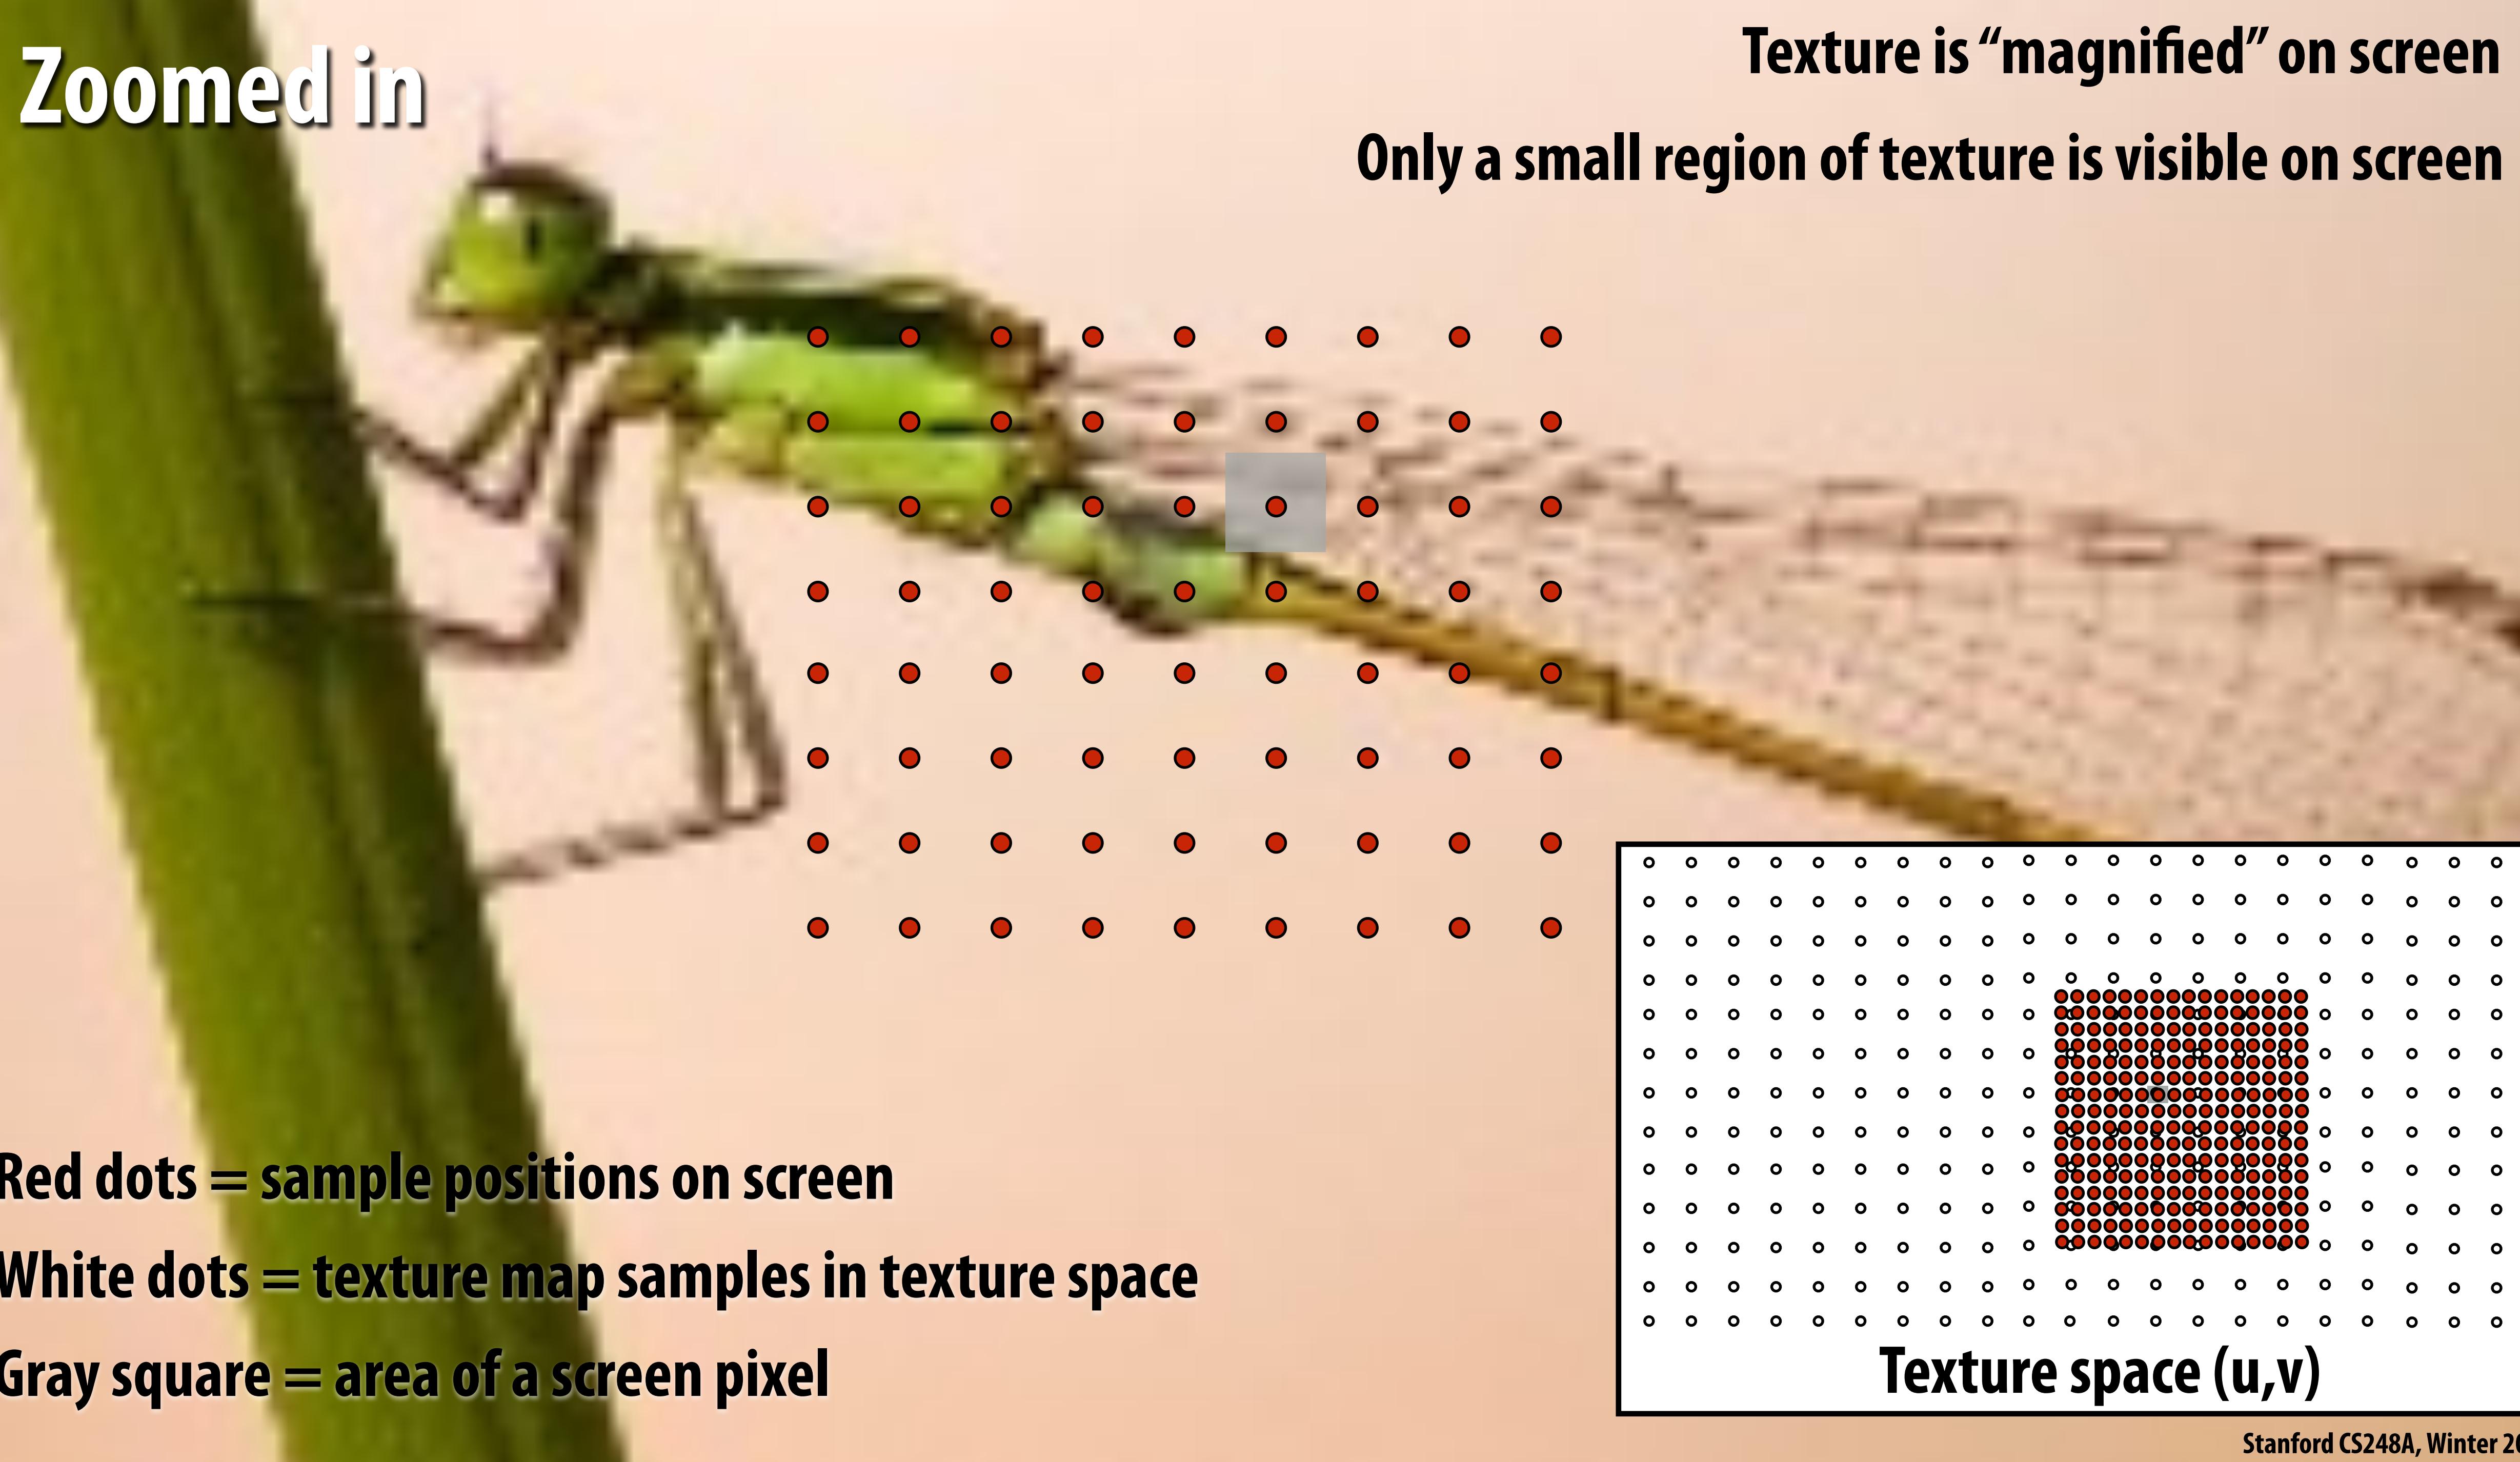
White dots = texture map samples in texture space

Gray square = area of a screen pixel

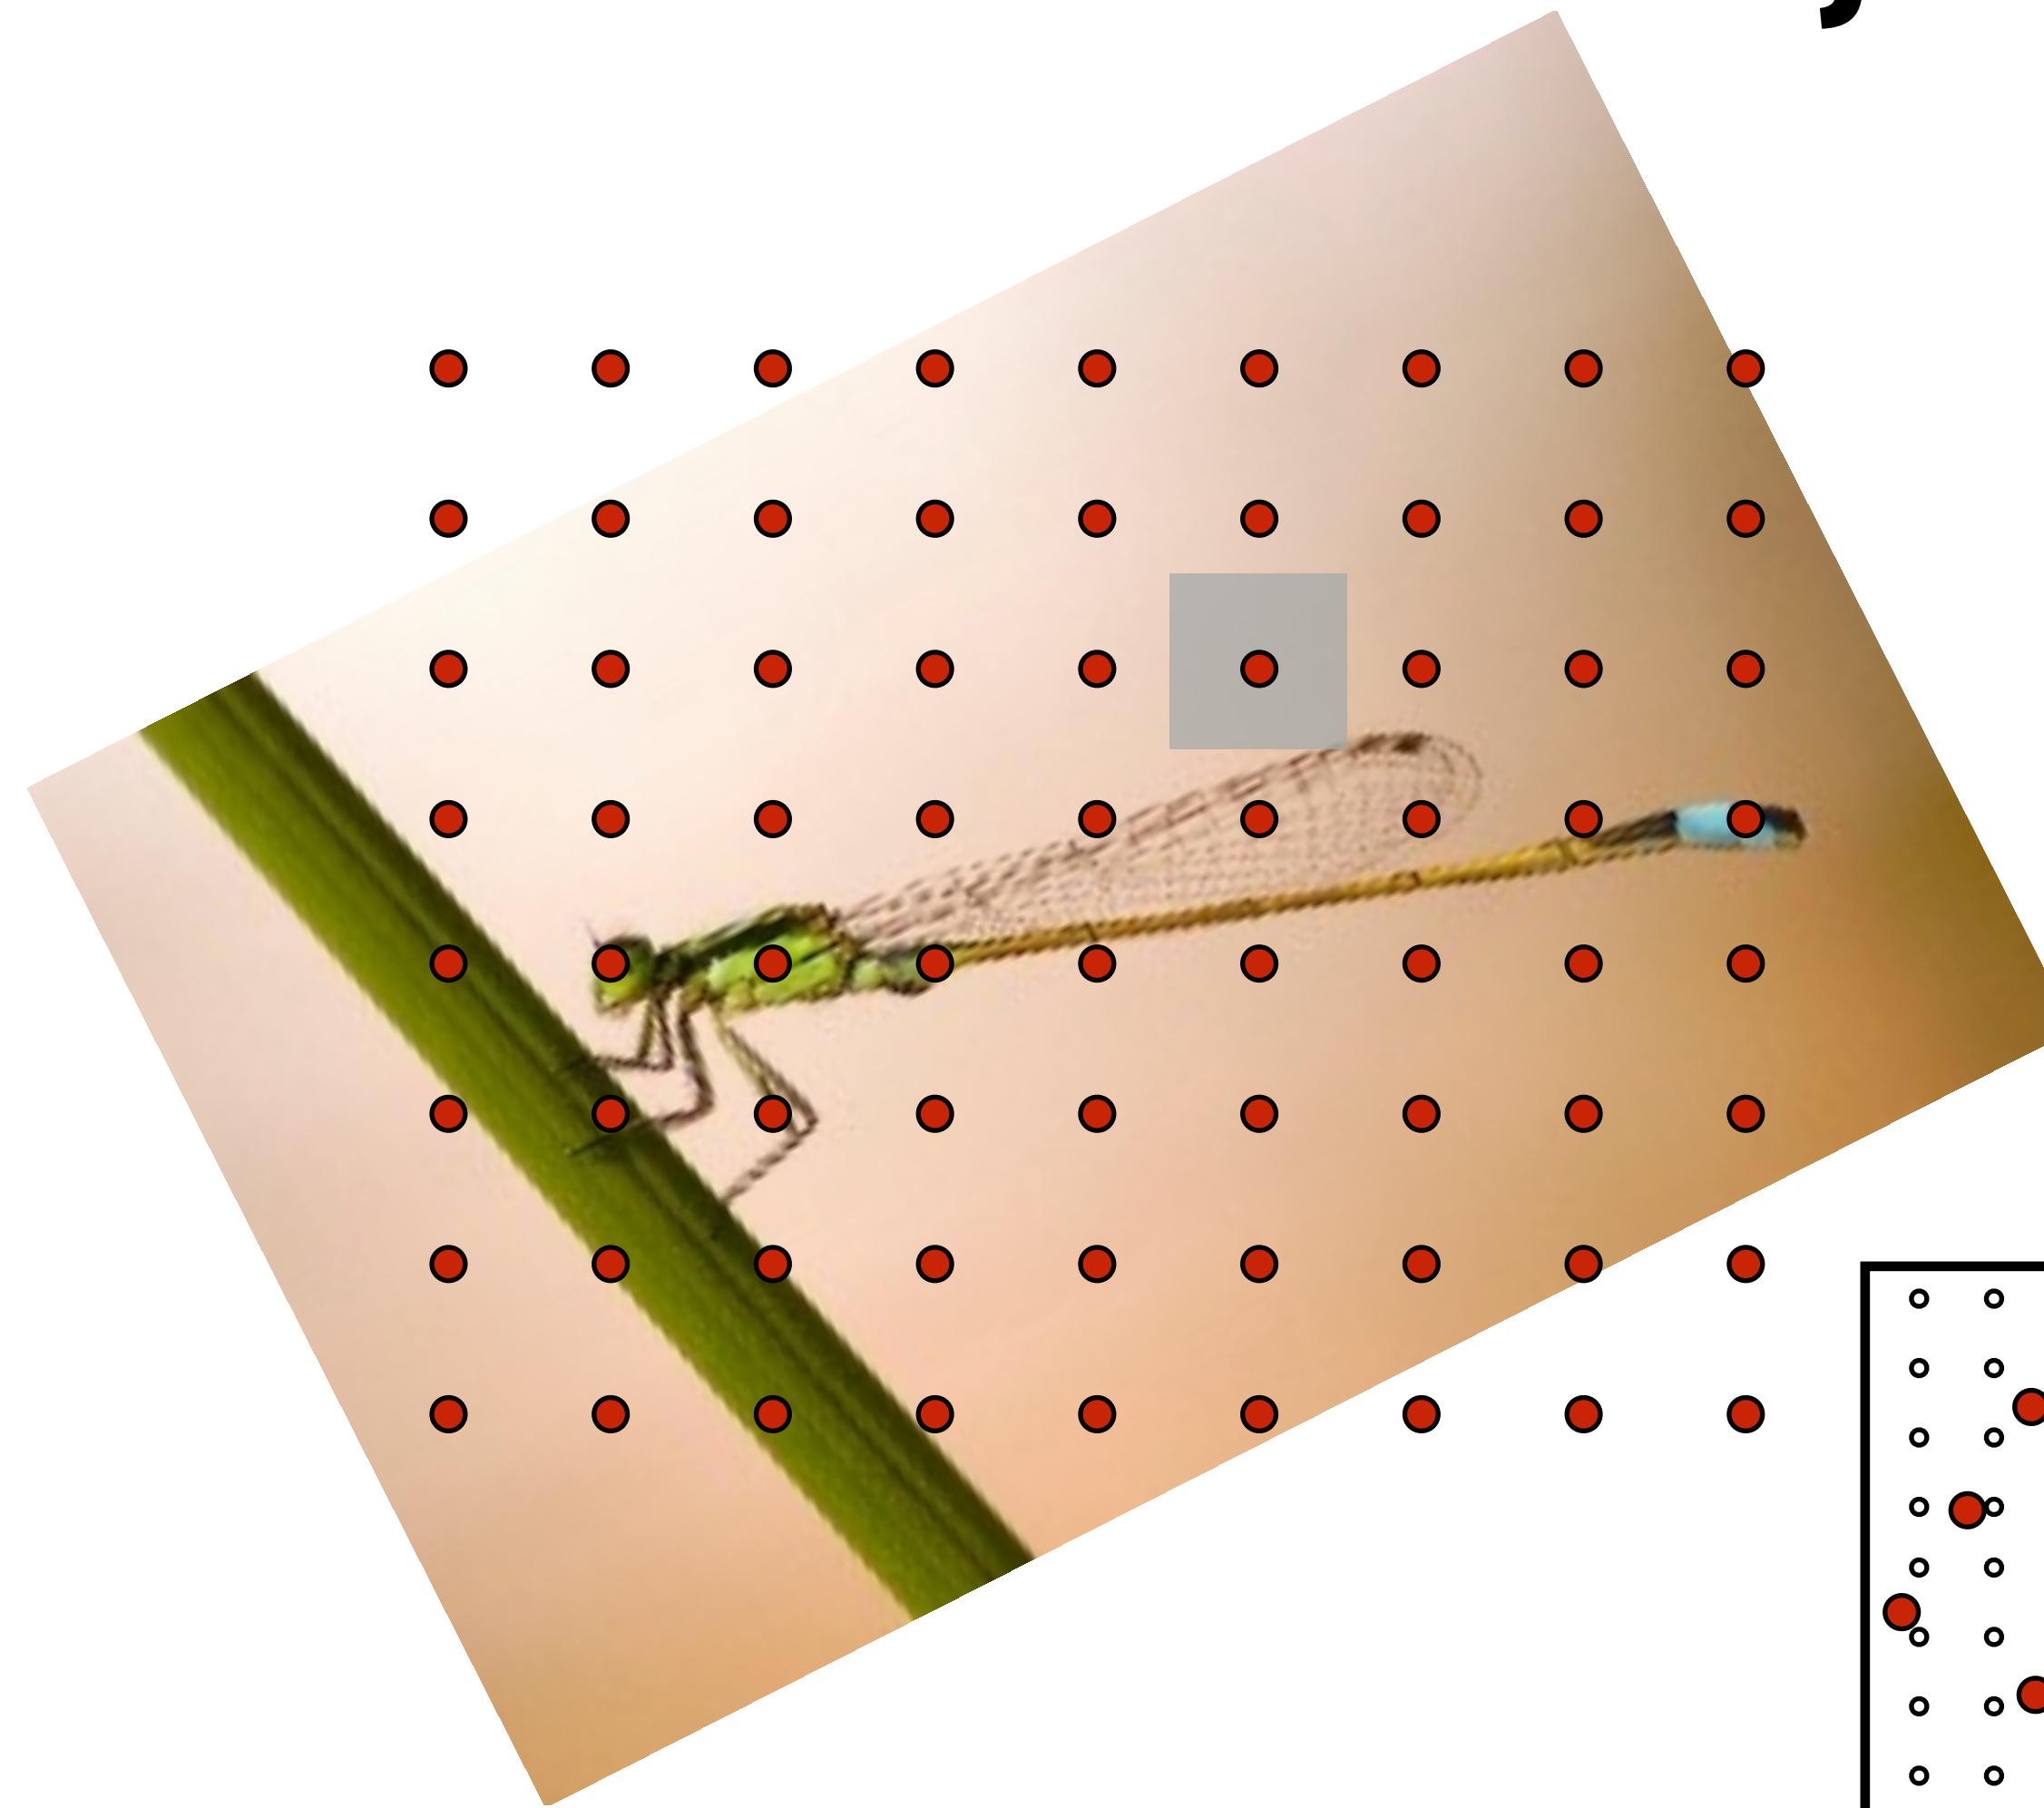


Texture space (u, v)

Zoomed in



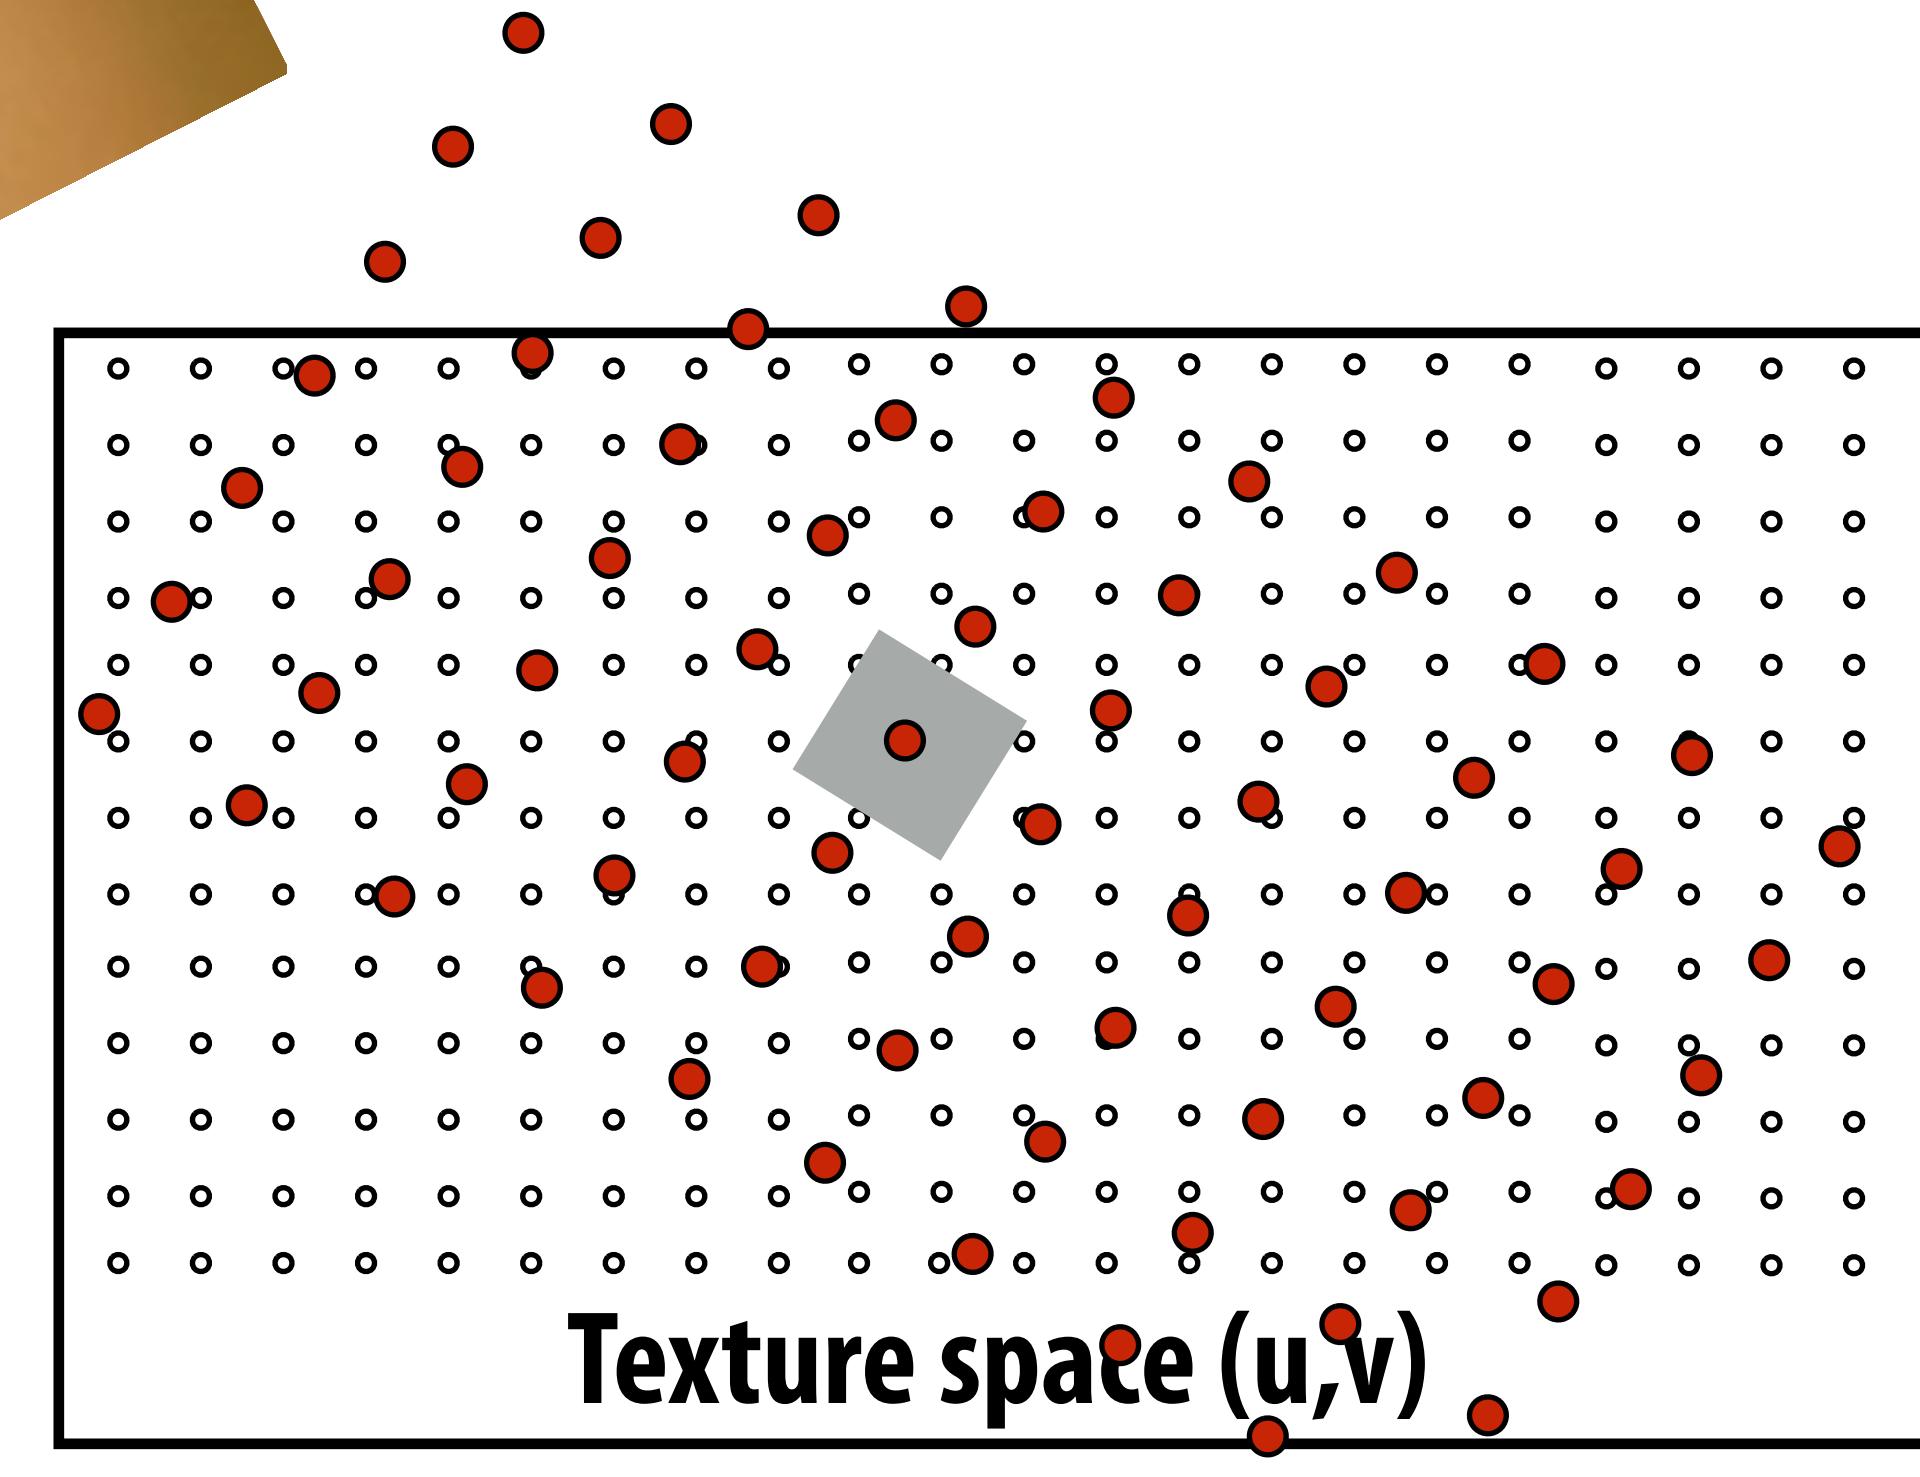
Sampling rate on screen vs in texture: object rotation



Red dots = sample positions on screen

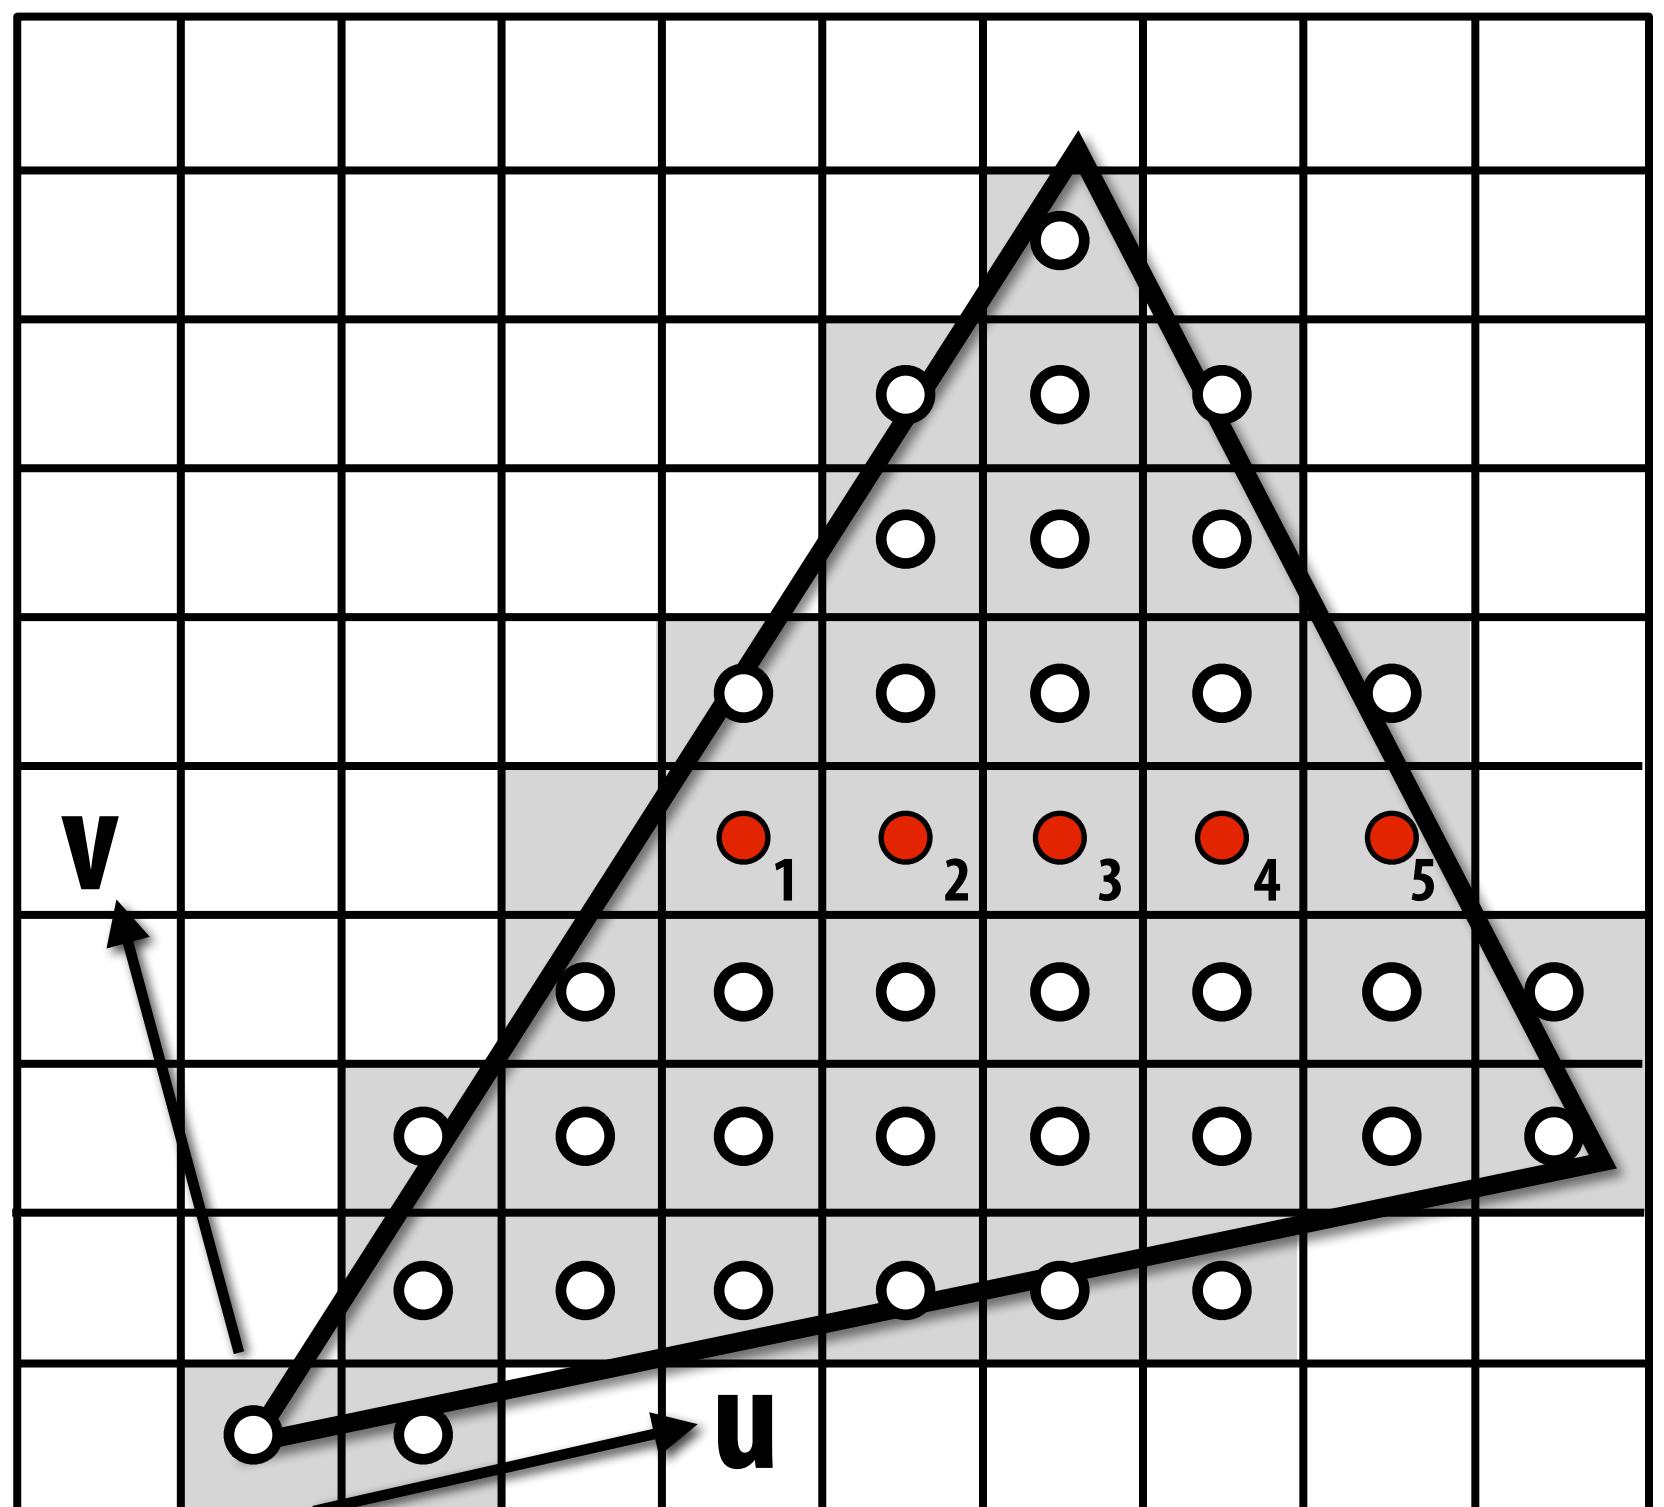
White dots = texture map samples in texture space

Gray square = area of a screen pixel



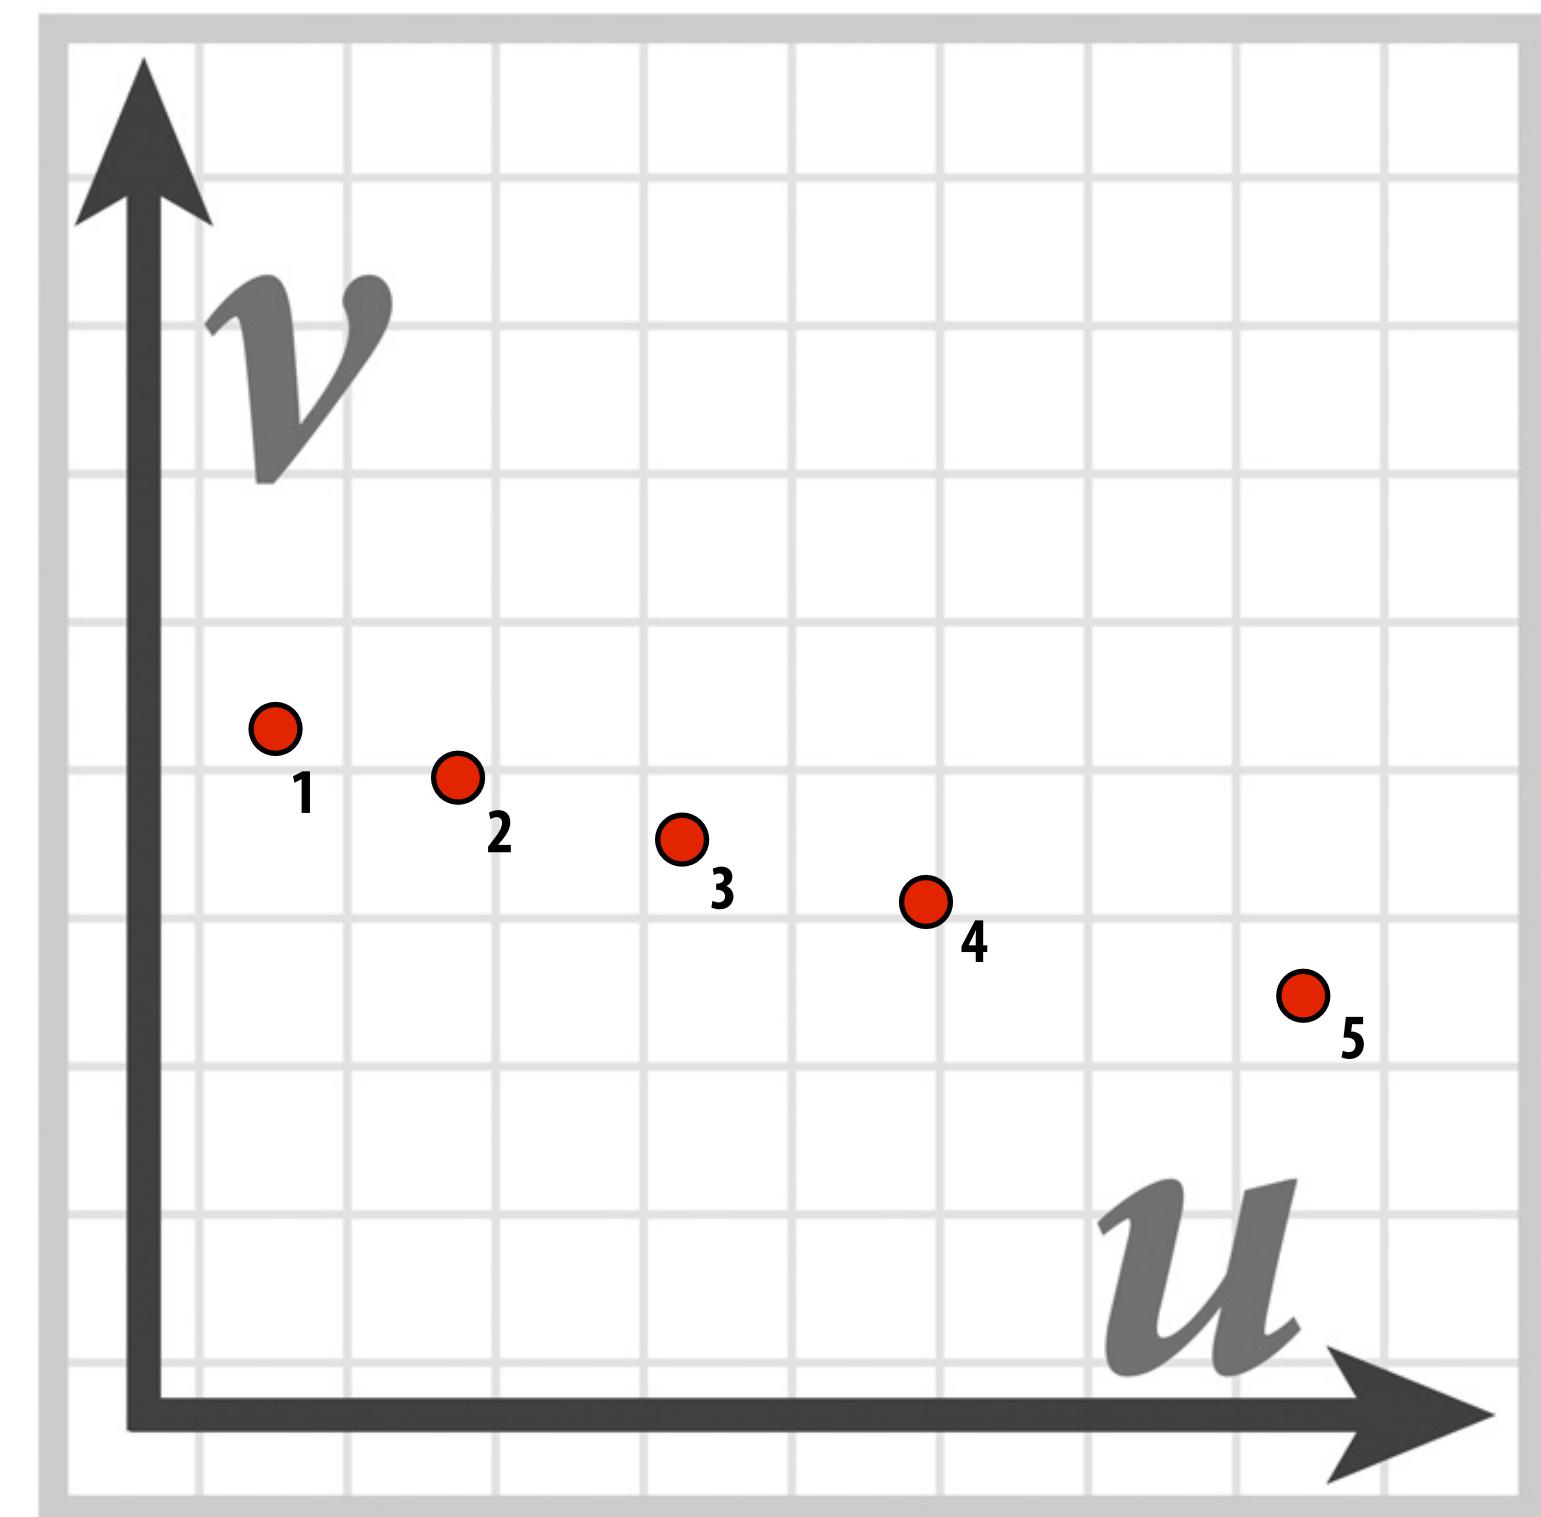
Equally spaced samples on screen != equally spaced samples in texture space

Sample positions in XY screen space



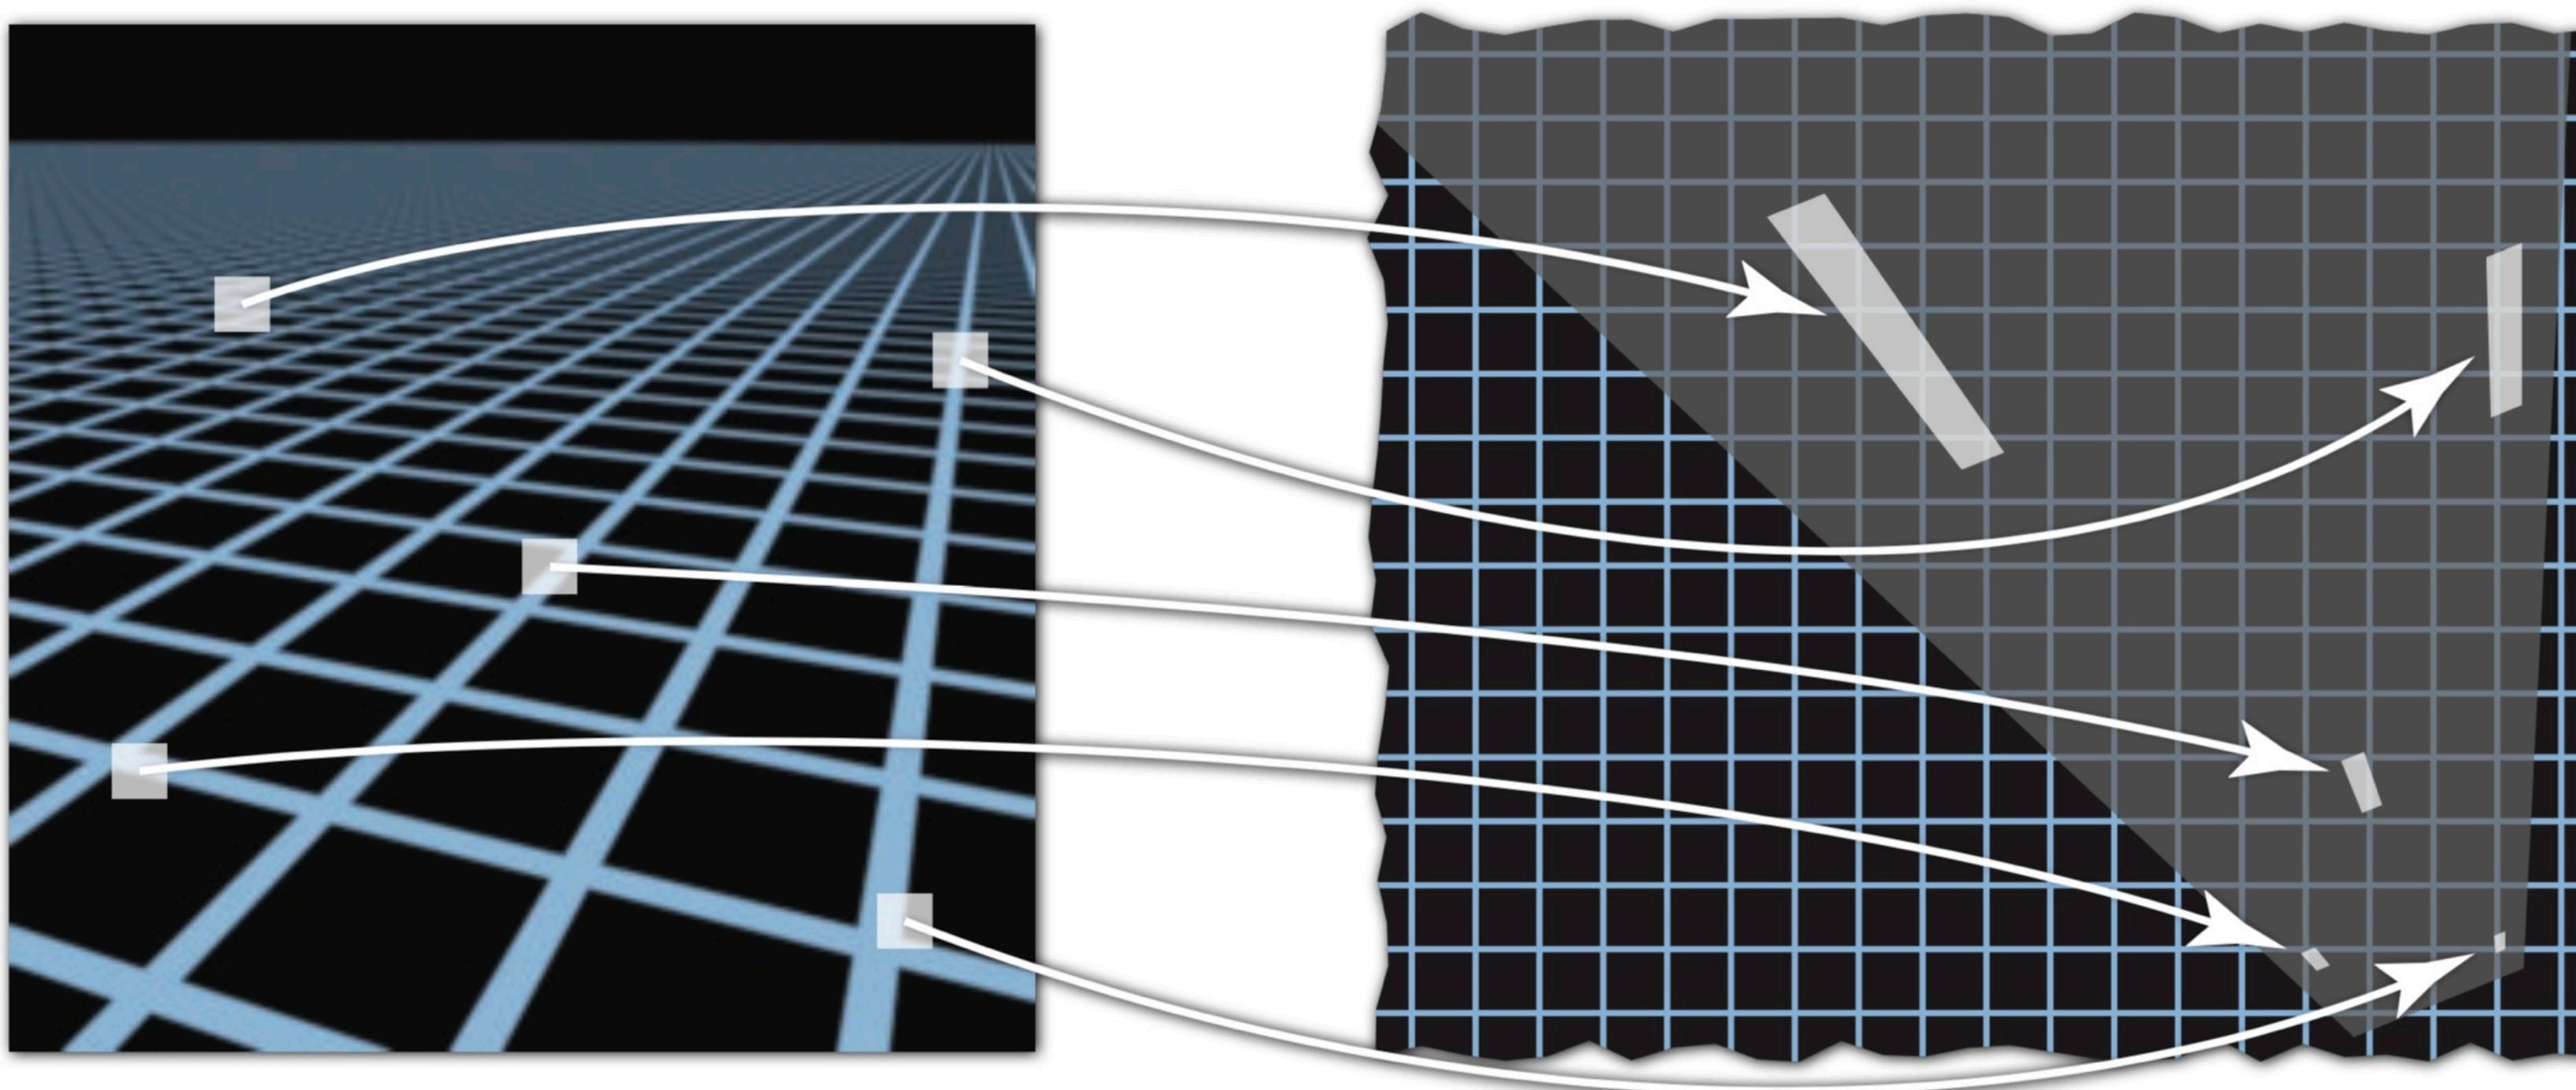
Sample positions are uniformly distributed in screen space
(renderer samples triangle's appearance at these locations)

Sample positions in texture space



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

Screen pixel footprint in texture space

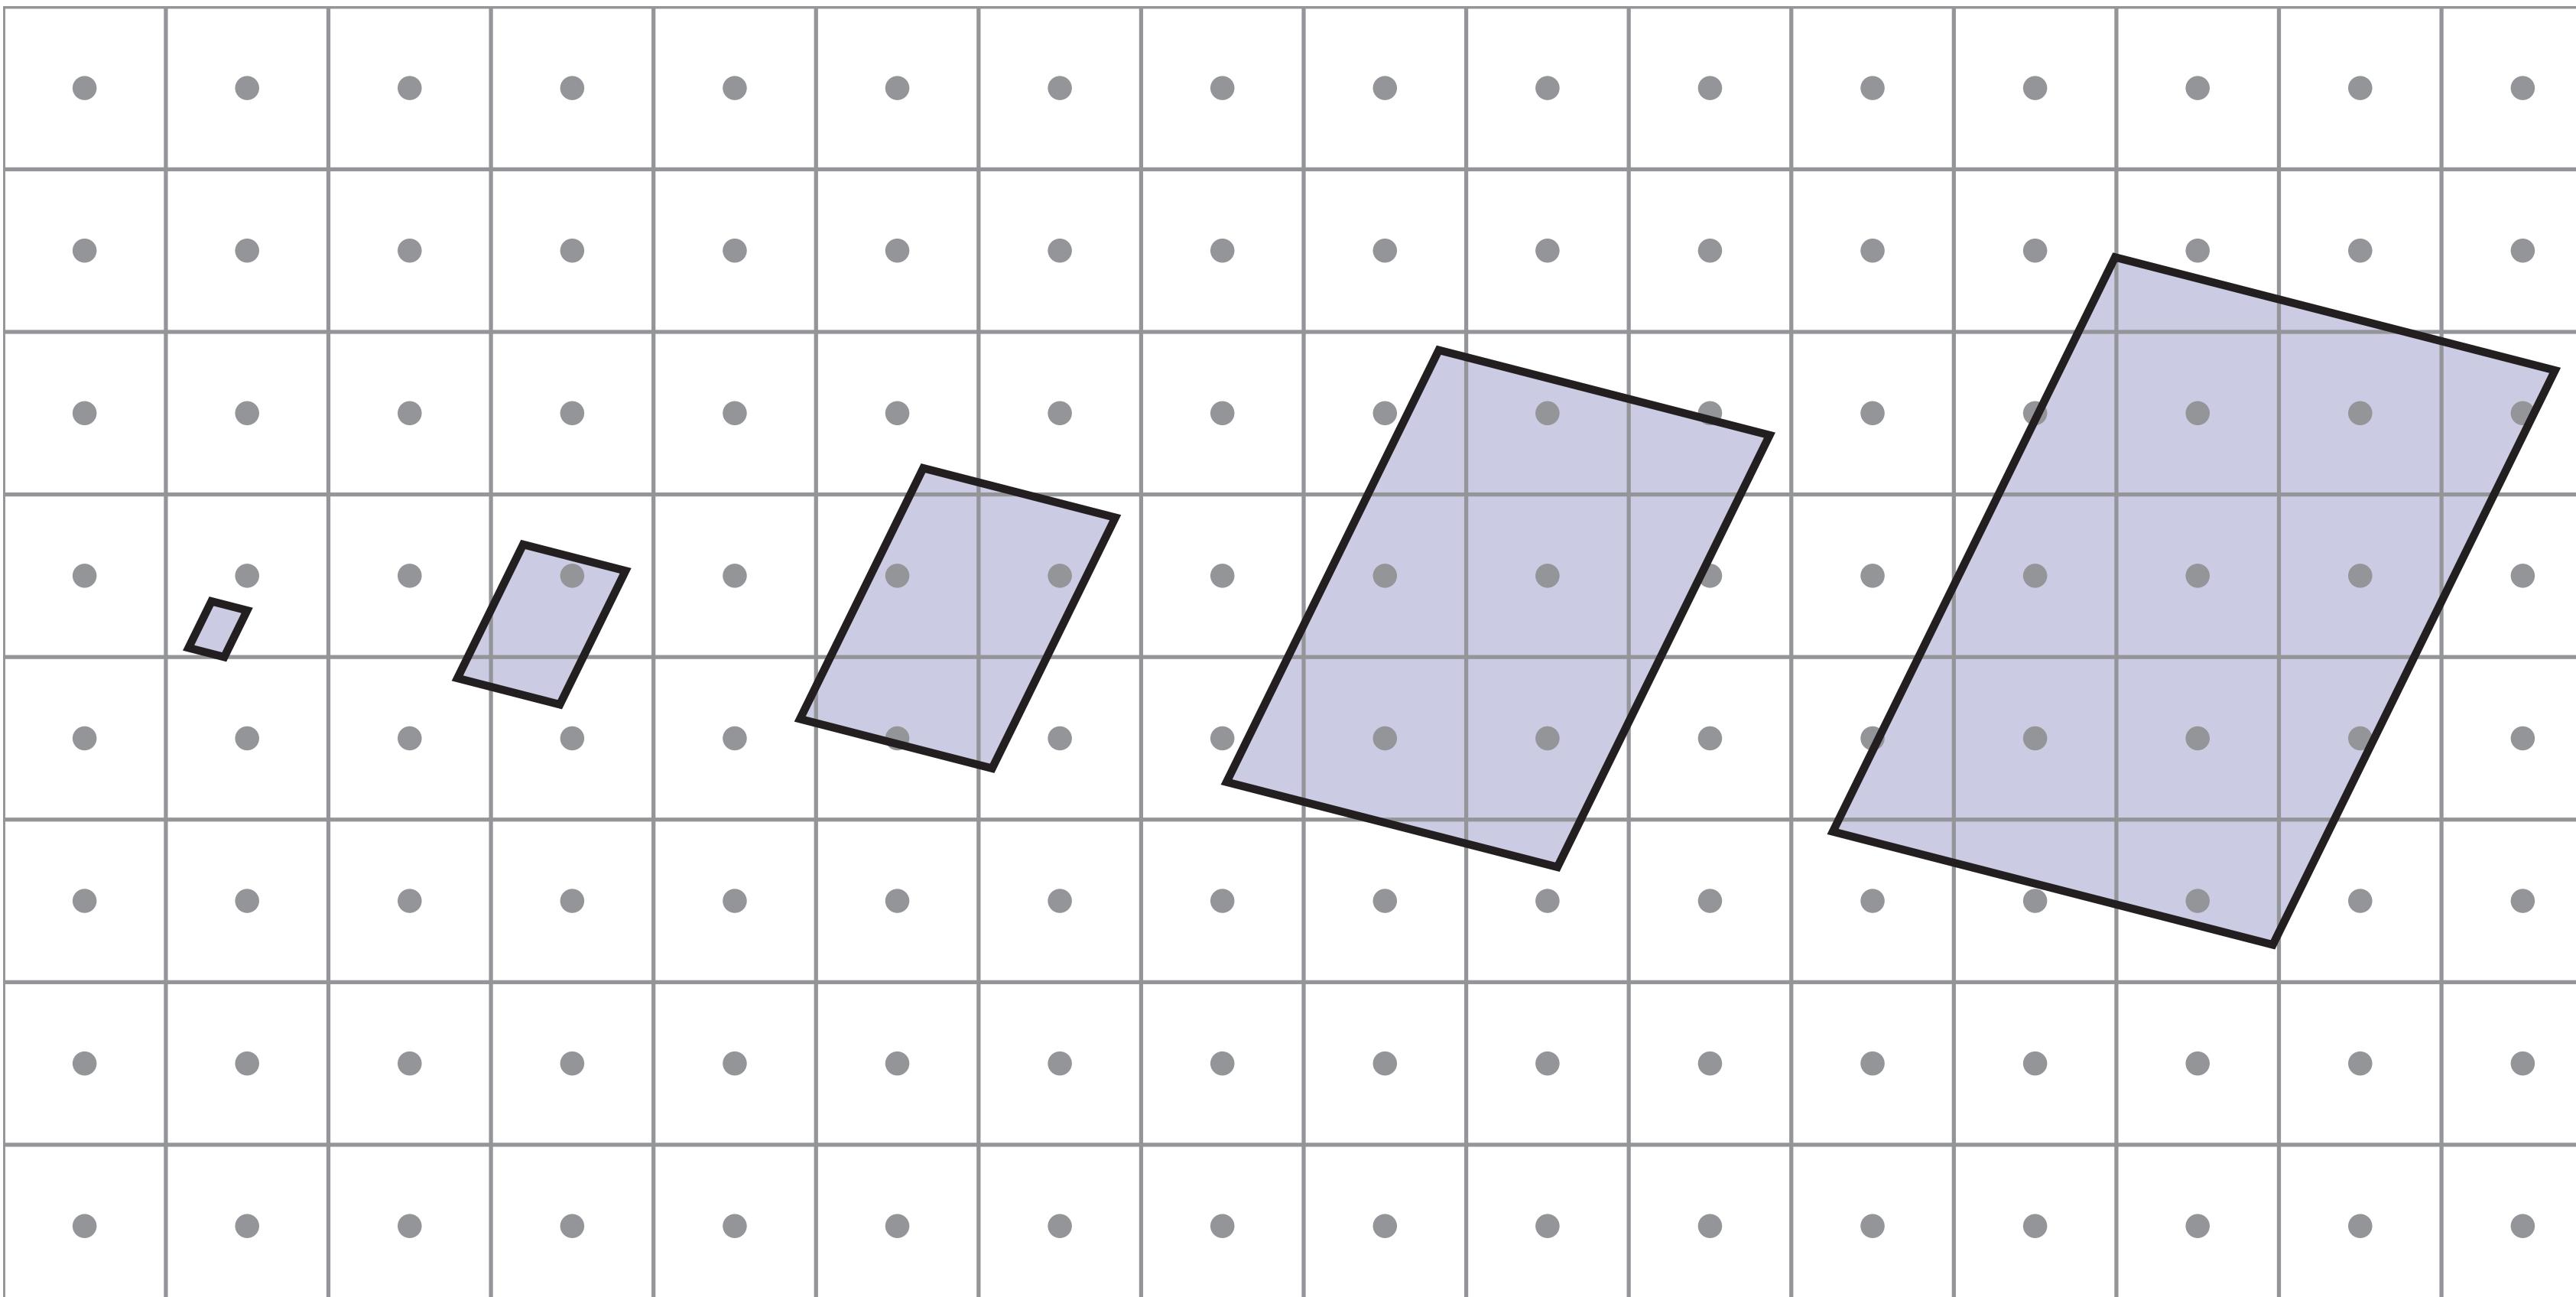


Screen space

Texture space

Texture sampling pattern not rectilinear or isotropic

Screen pixel footprint in texture space



**Upsampling
(Magnification)**

*Camera zoomed in
close to object*

**Downsampling
(Minification)**

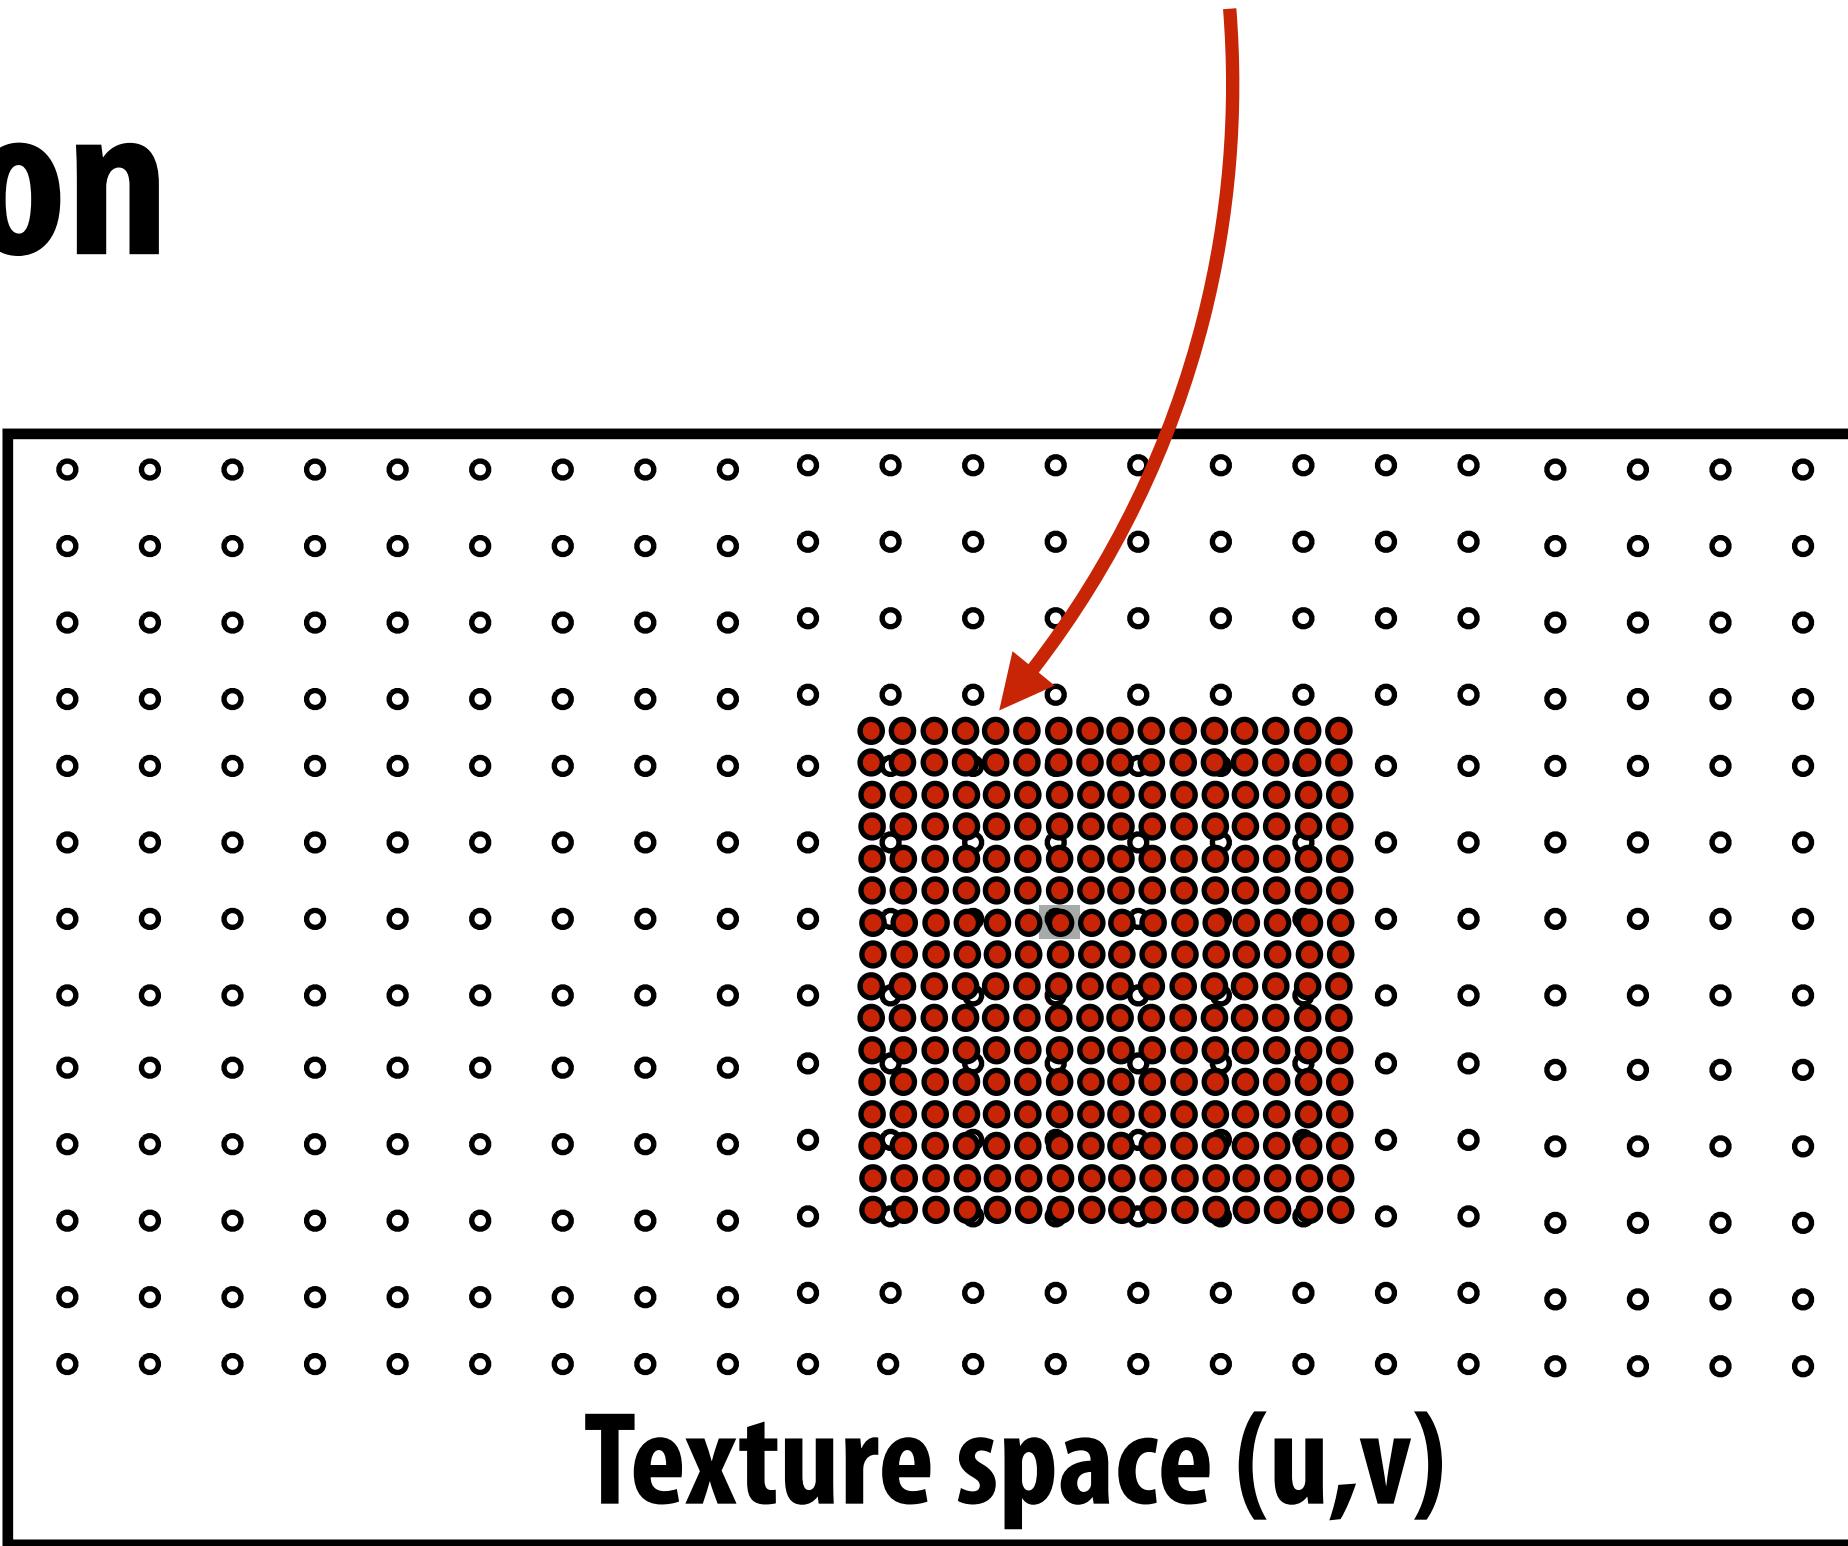
*Camera far away
from object*

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!
- When pixel area is larger than texel area (texture minification)
 - Think: zoom far out from object
 - One pixel sample per multiple texel samples
- When pixel area is smaller than texel area (texture magnification)
 - Think: zoom in on an object
 - Multiple pixel samples per texel sample

Texture magnification

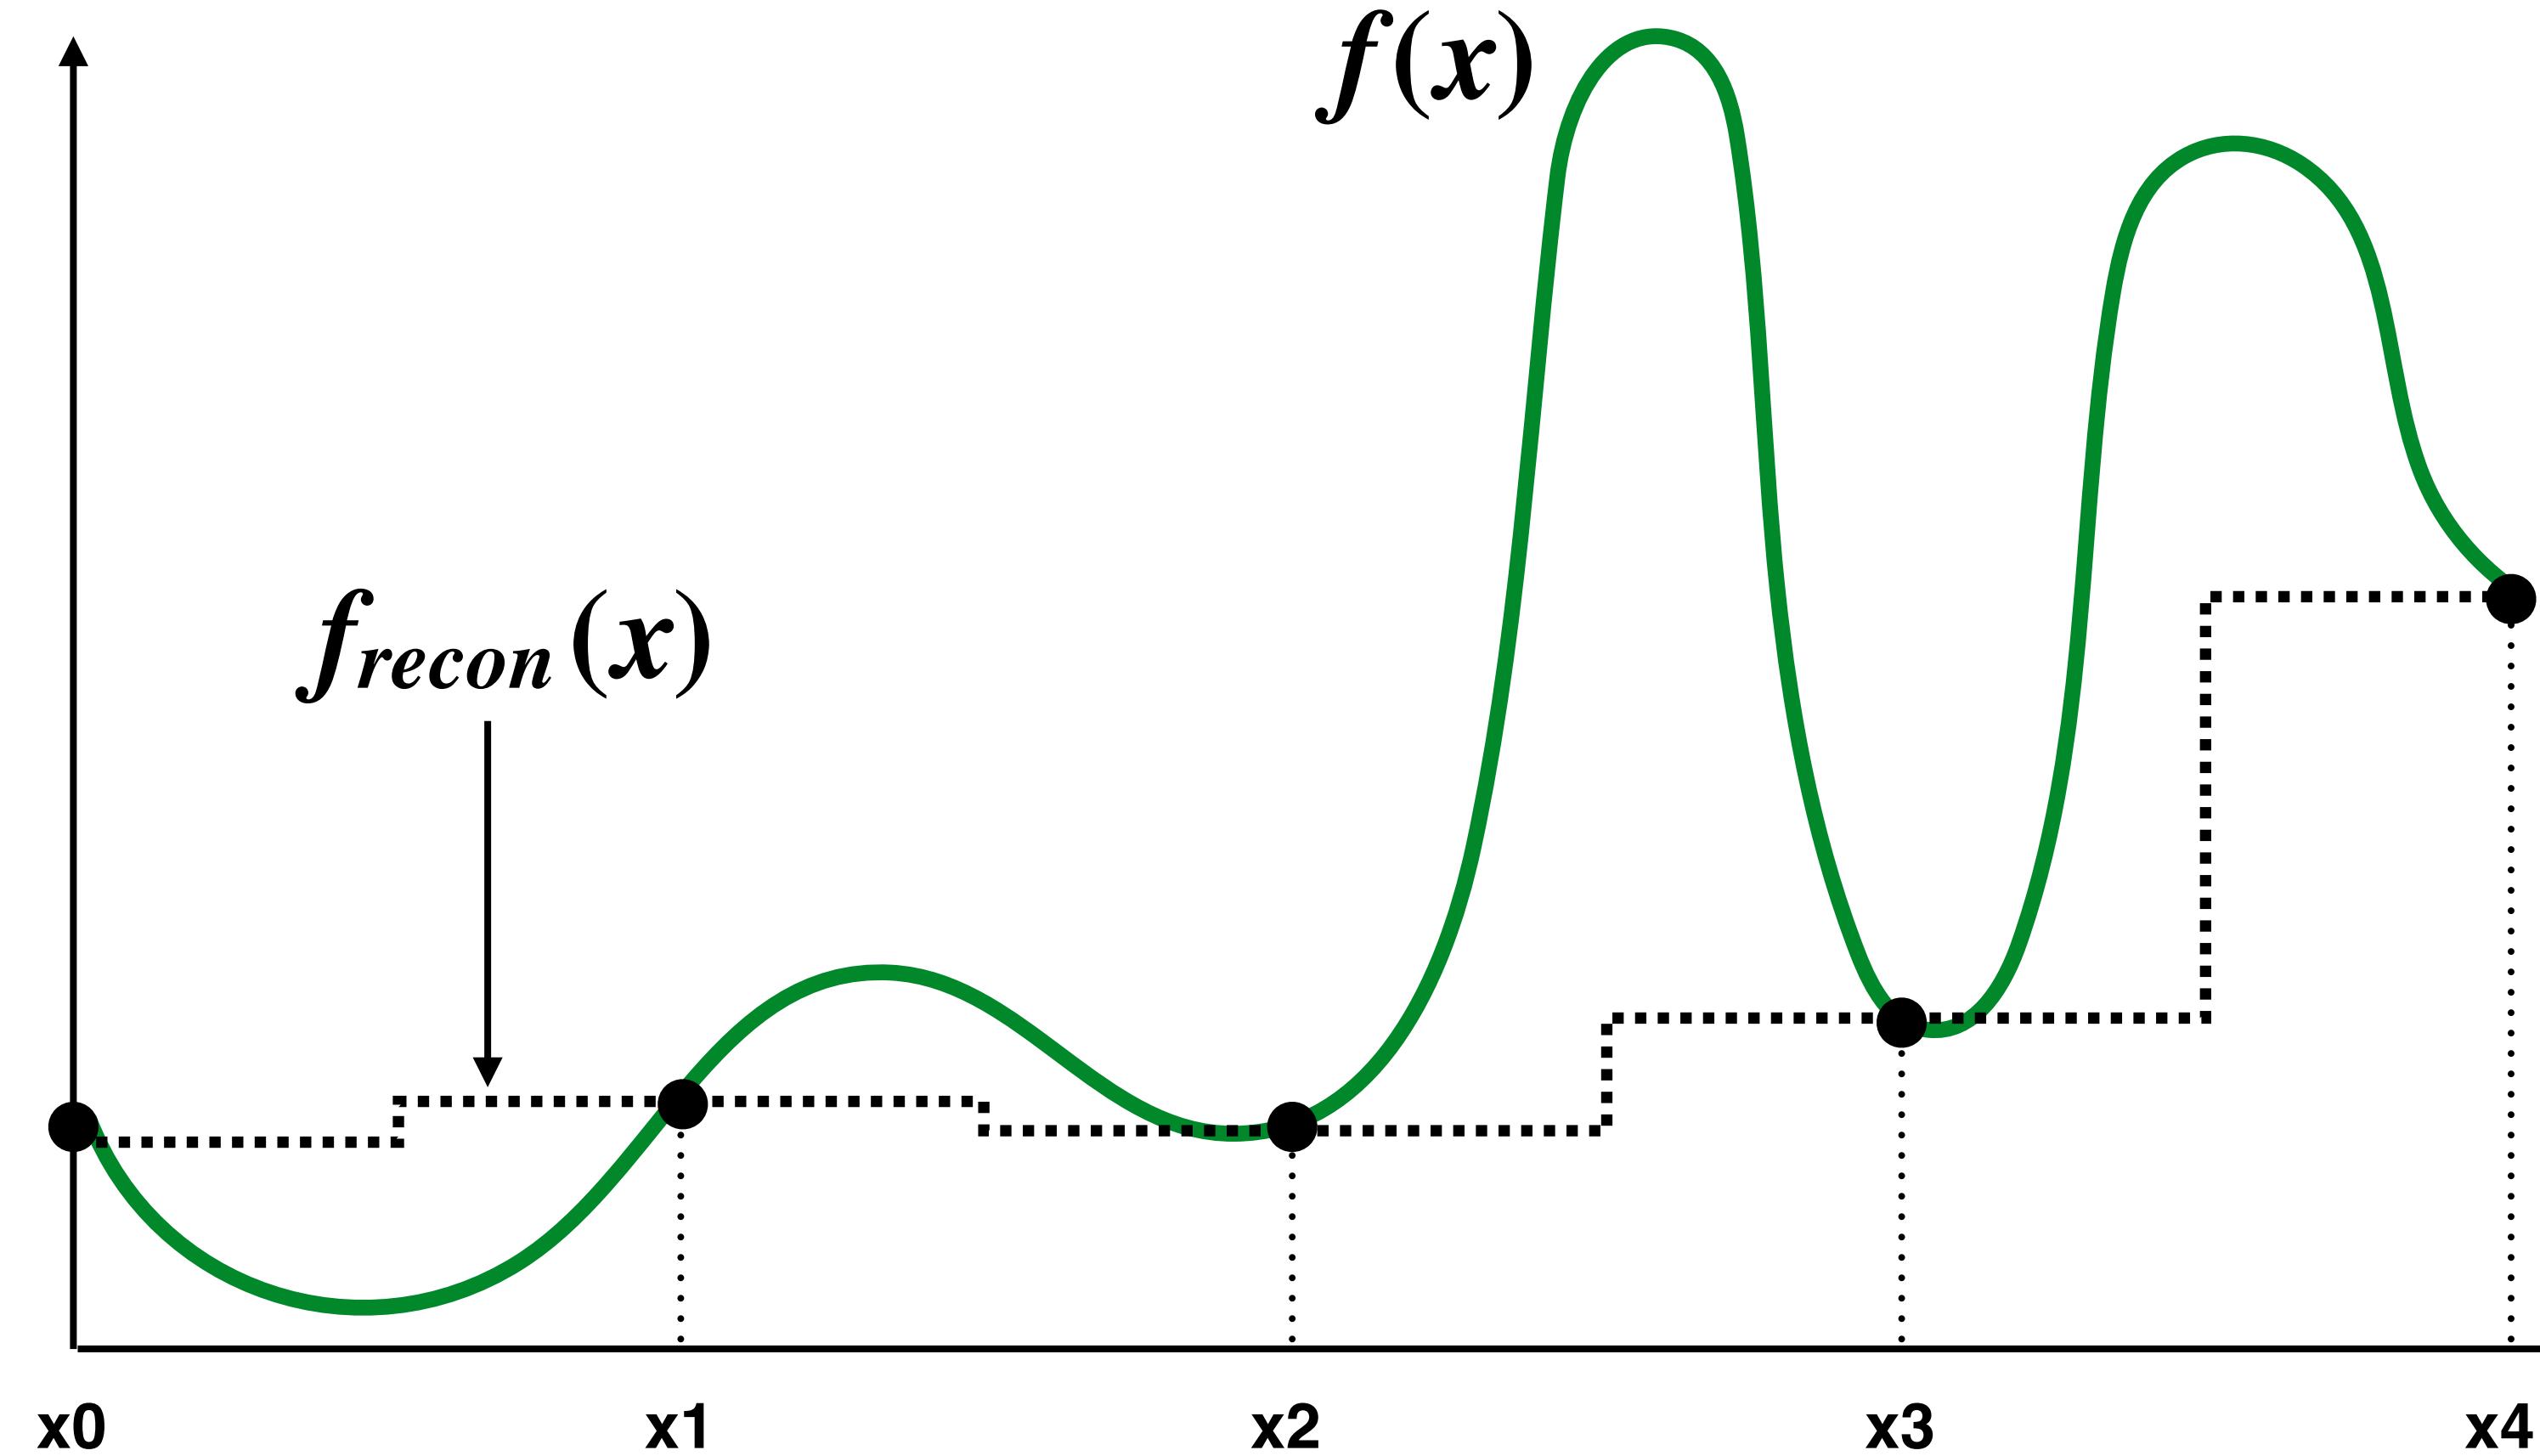
What is the color of the texture
at these red dots?



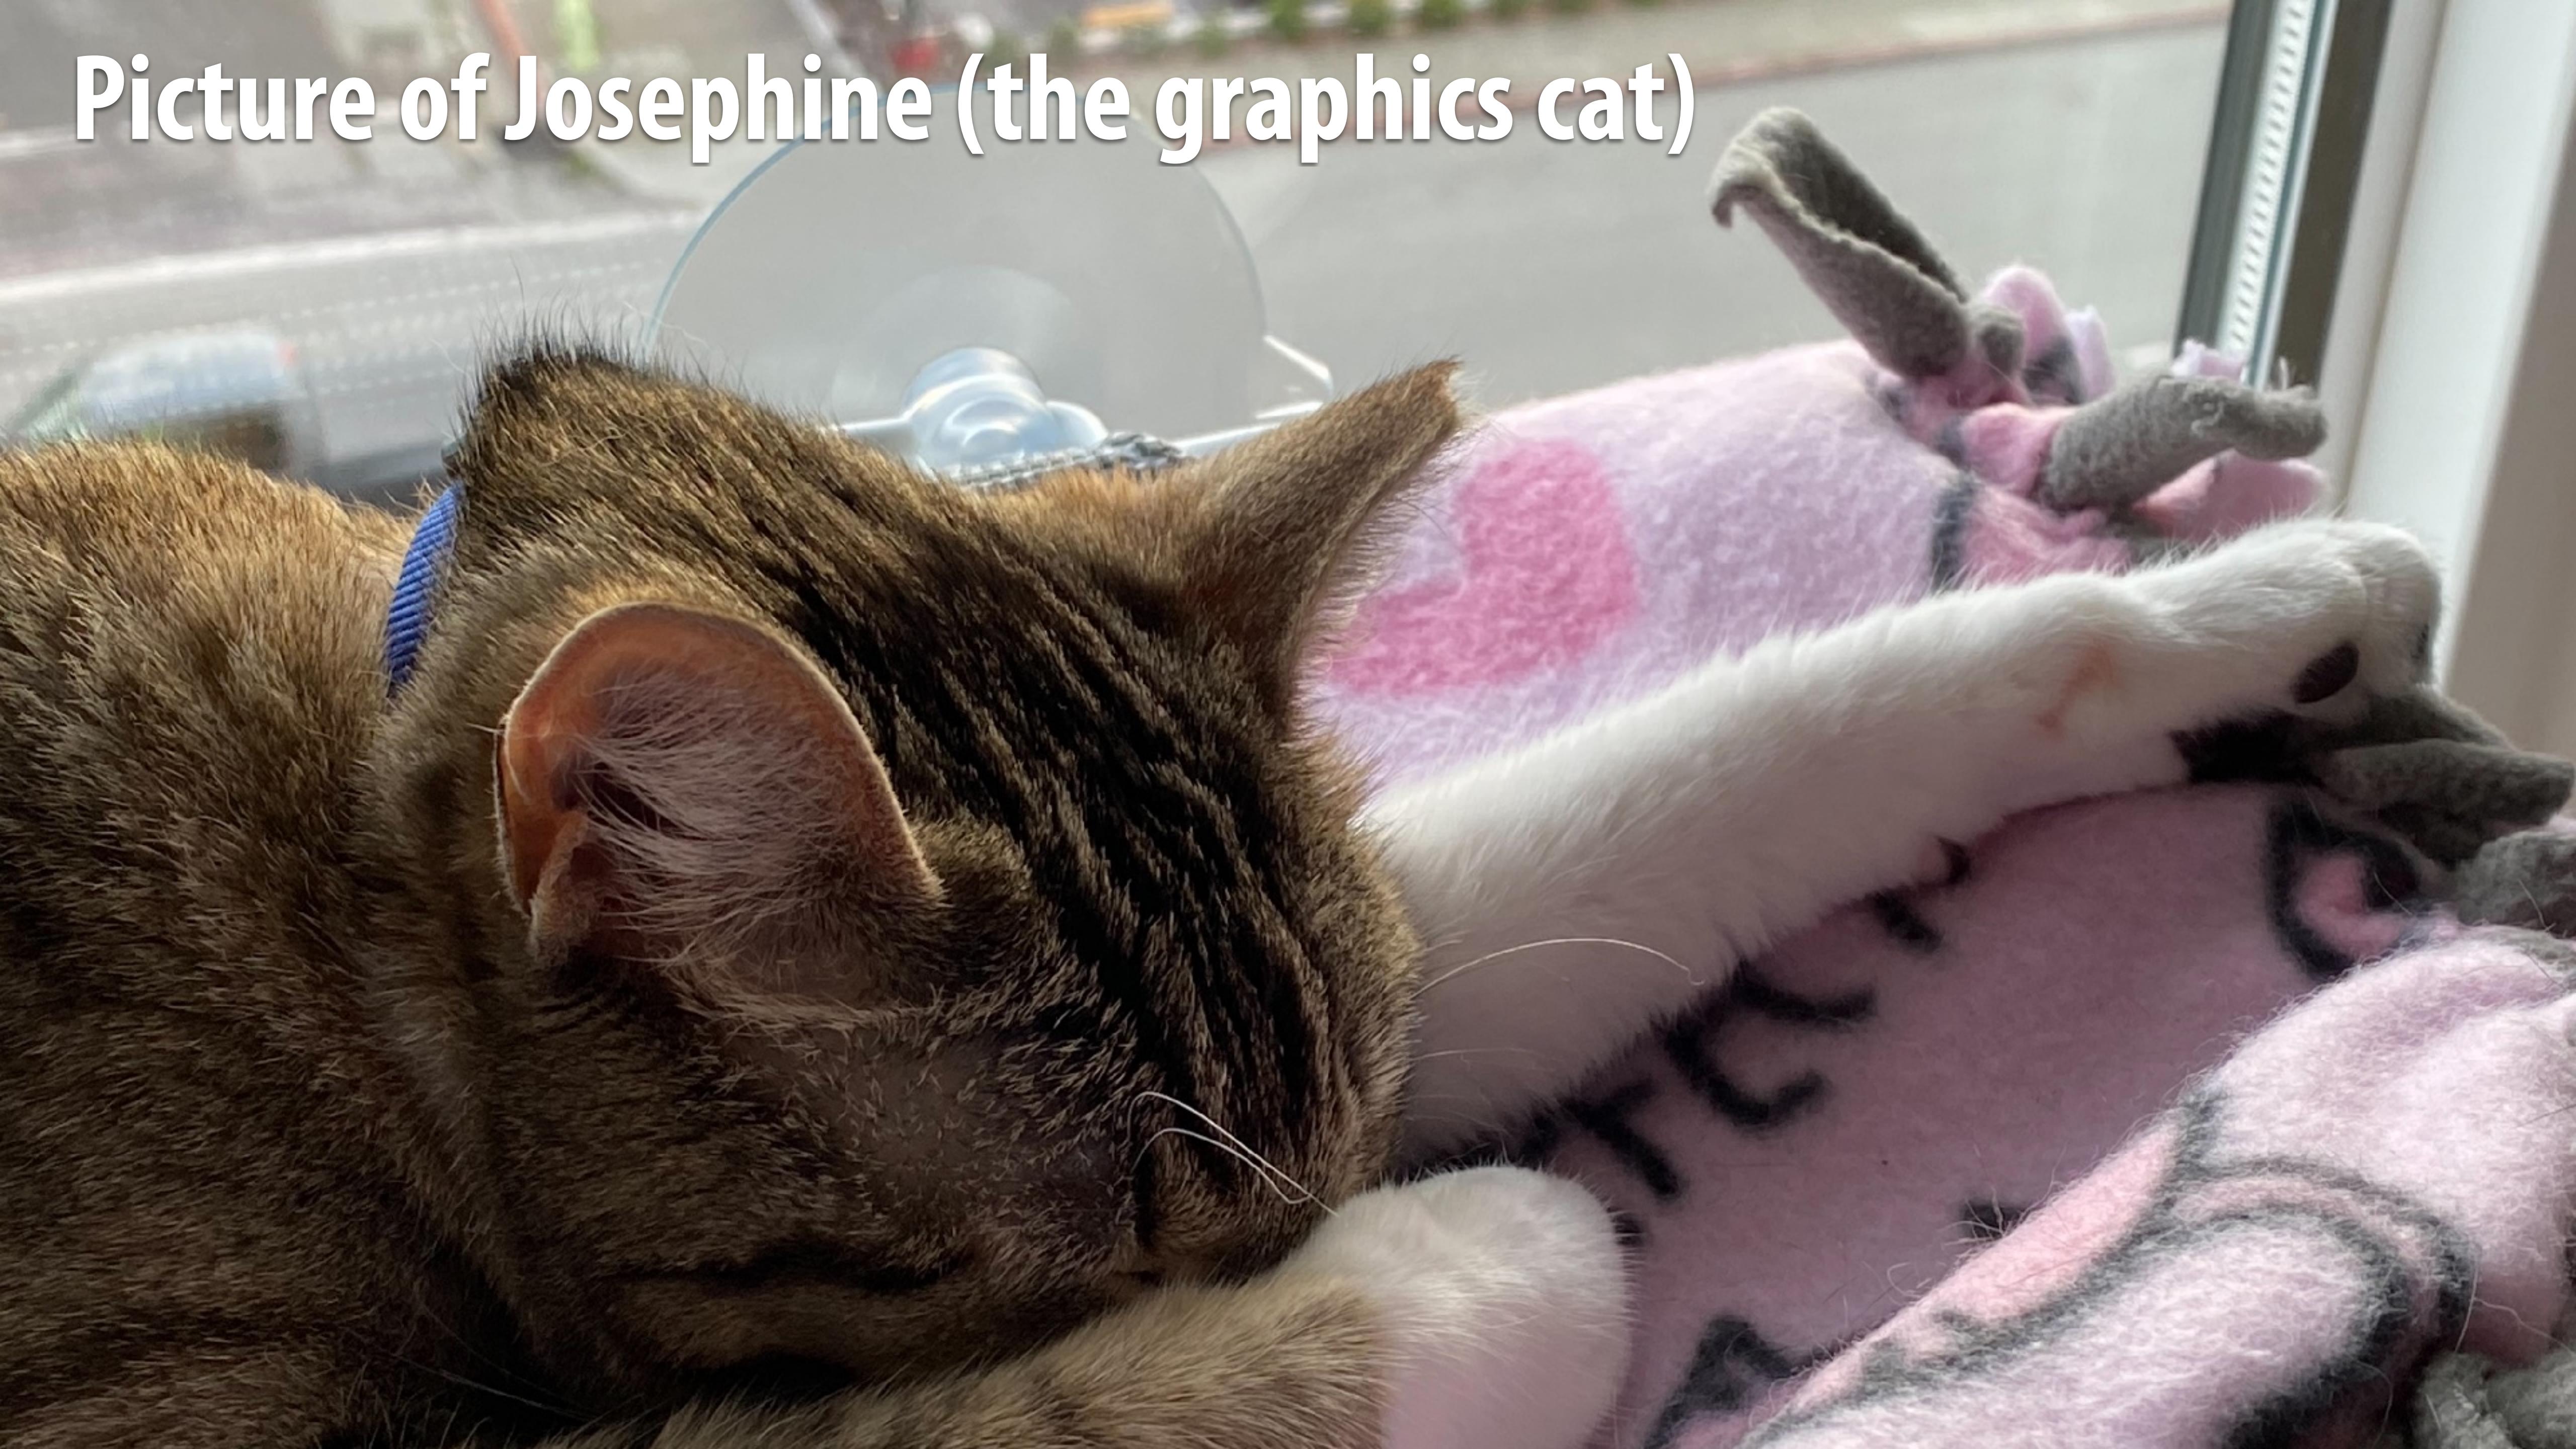
Review: piecewise constant approximation

$f_{recon}(x)$ = value of sample closest to x

$f_{recon}(x)$ approximates $f(x)$



Picture of Josephine (the graphics cat)



Texture magnification (nearest)



Texture magnification (nearest)



Texture magnification (nearest)



Texture magnification (nearest)



Texture magnification

- Generally don't want this situation — it means we have insufficient texture resolution
- Magnification involves interpolation of values in texture map (below: three different interpolation kernel functions)



Nearest sample



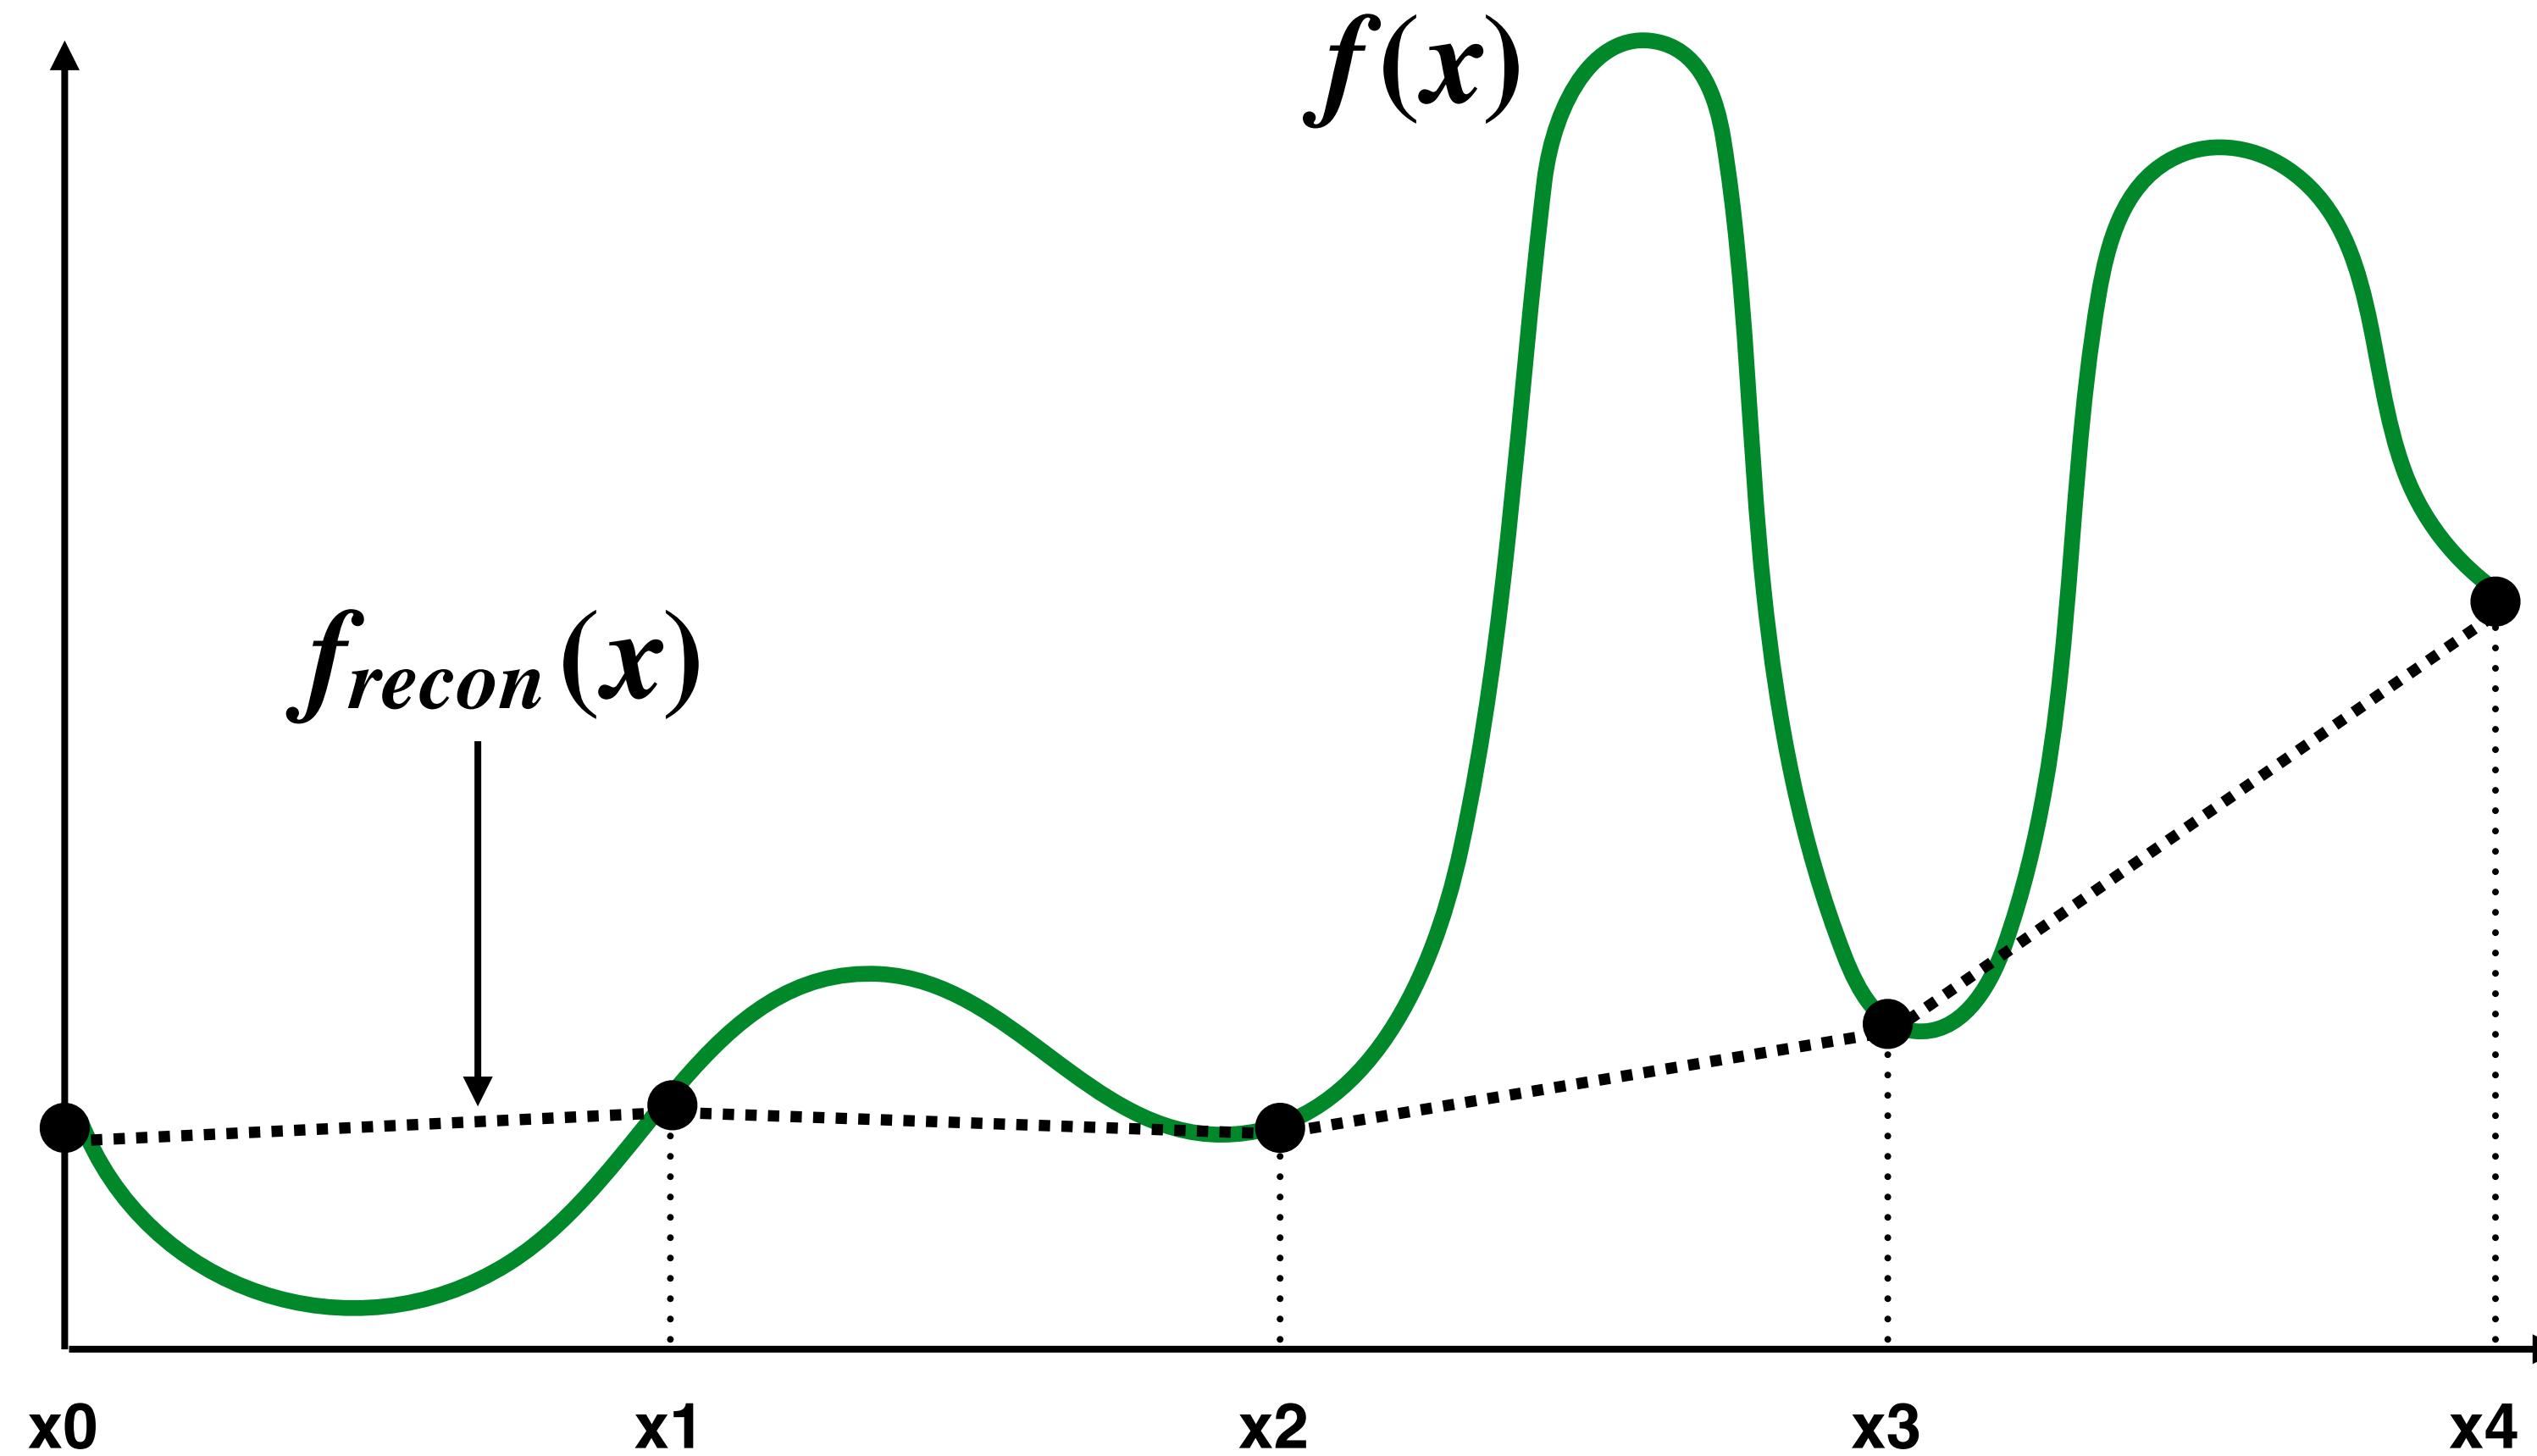
Bilinear



Bicubic

Review: piecewise linear approximation

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



Texture magnification (bilinear)



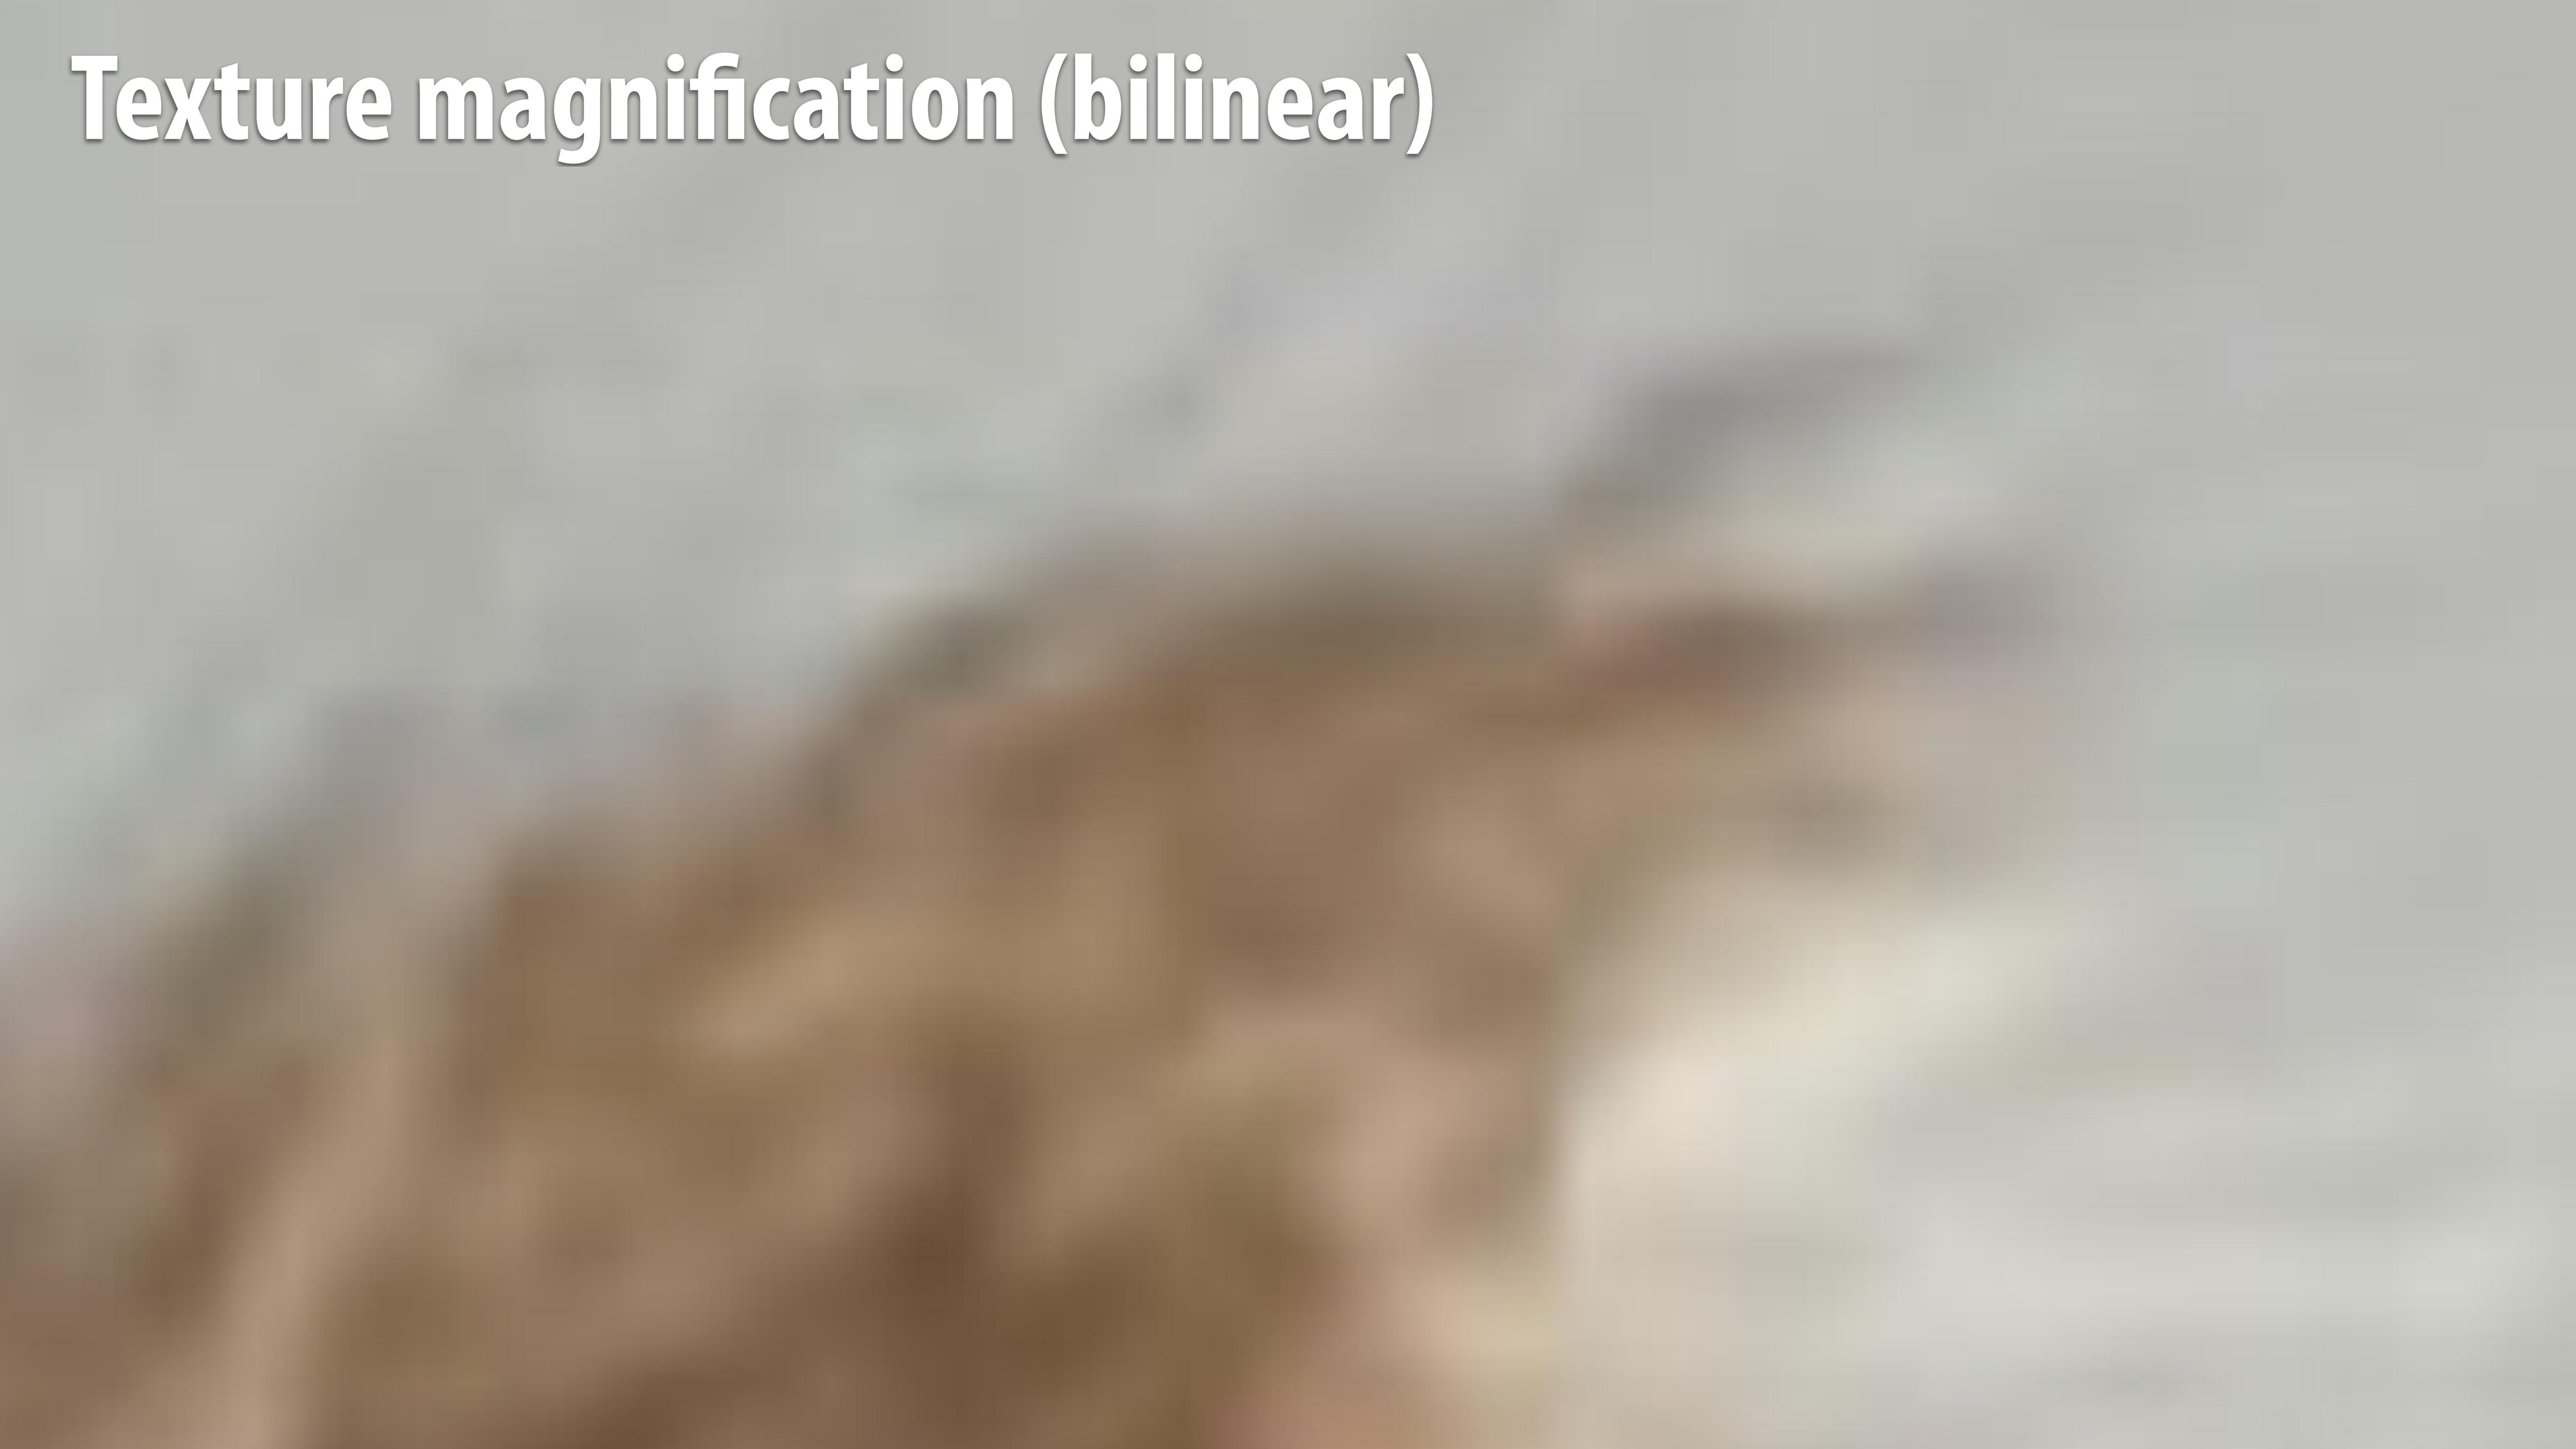
Texture magnification (bilinear)



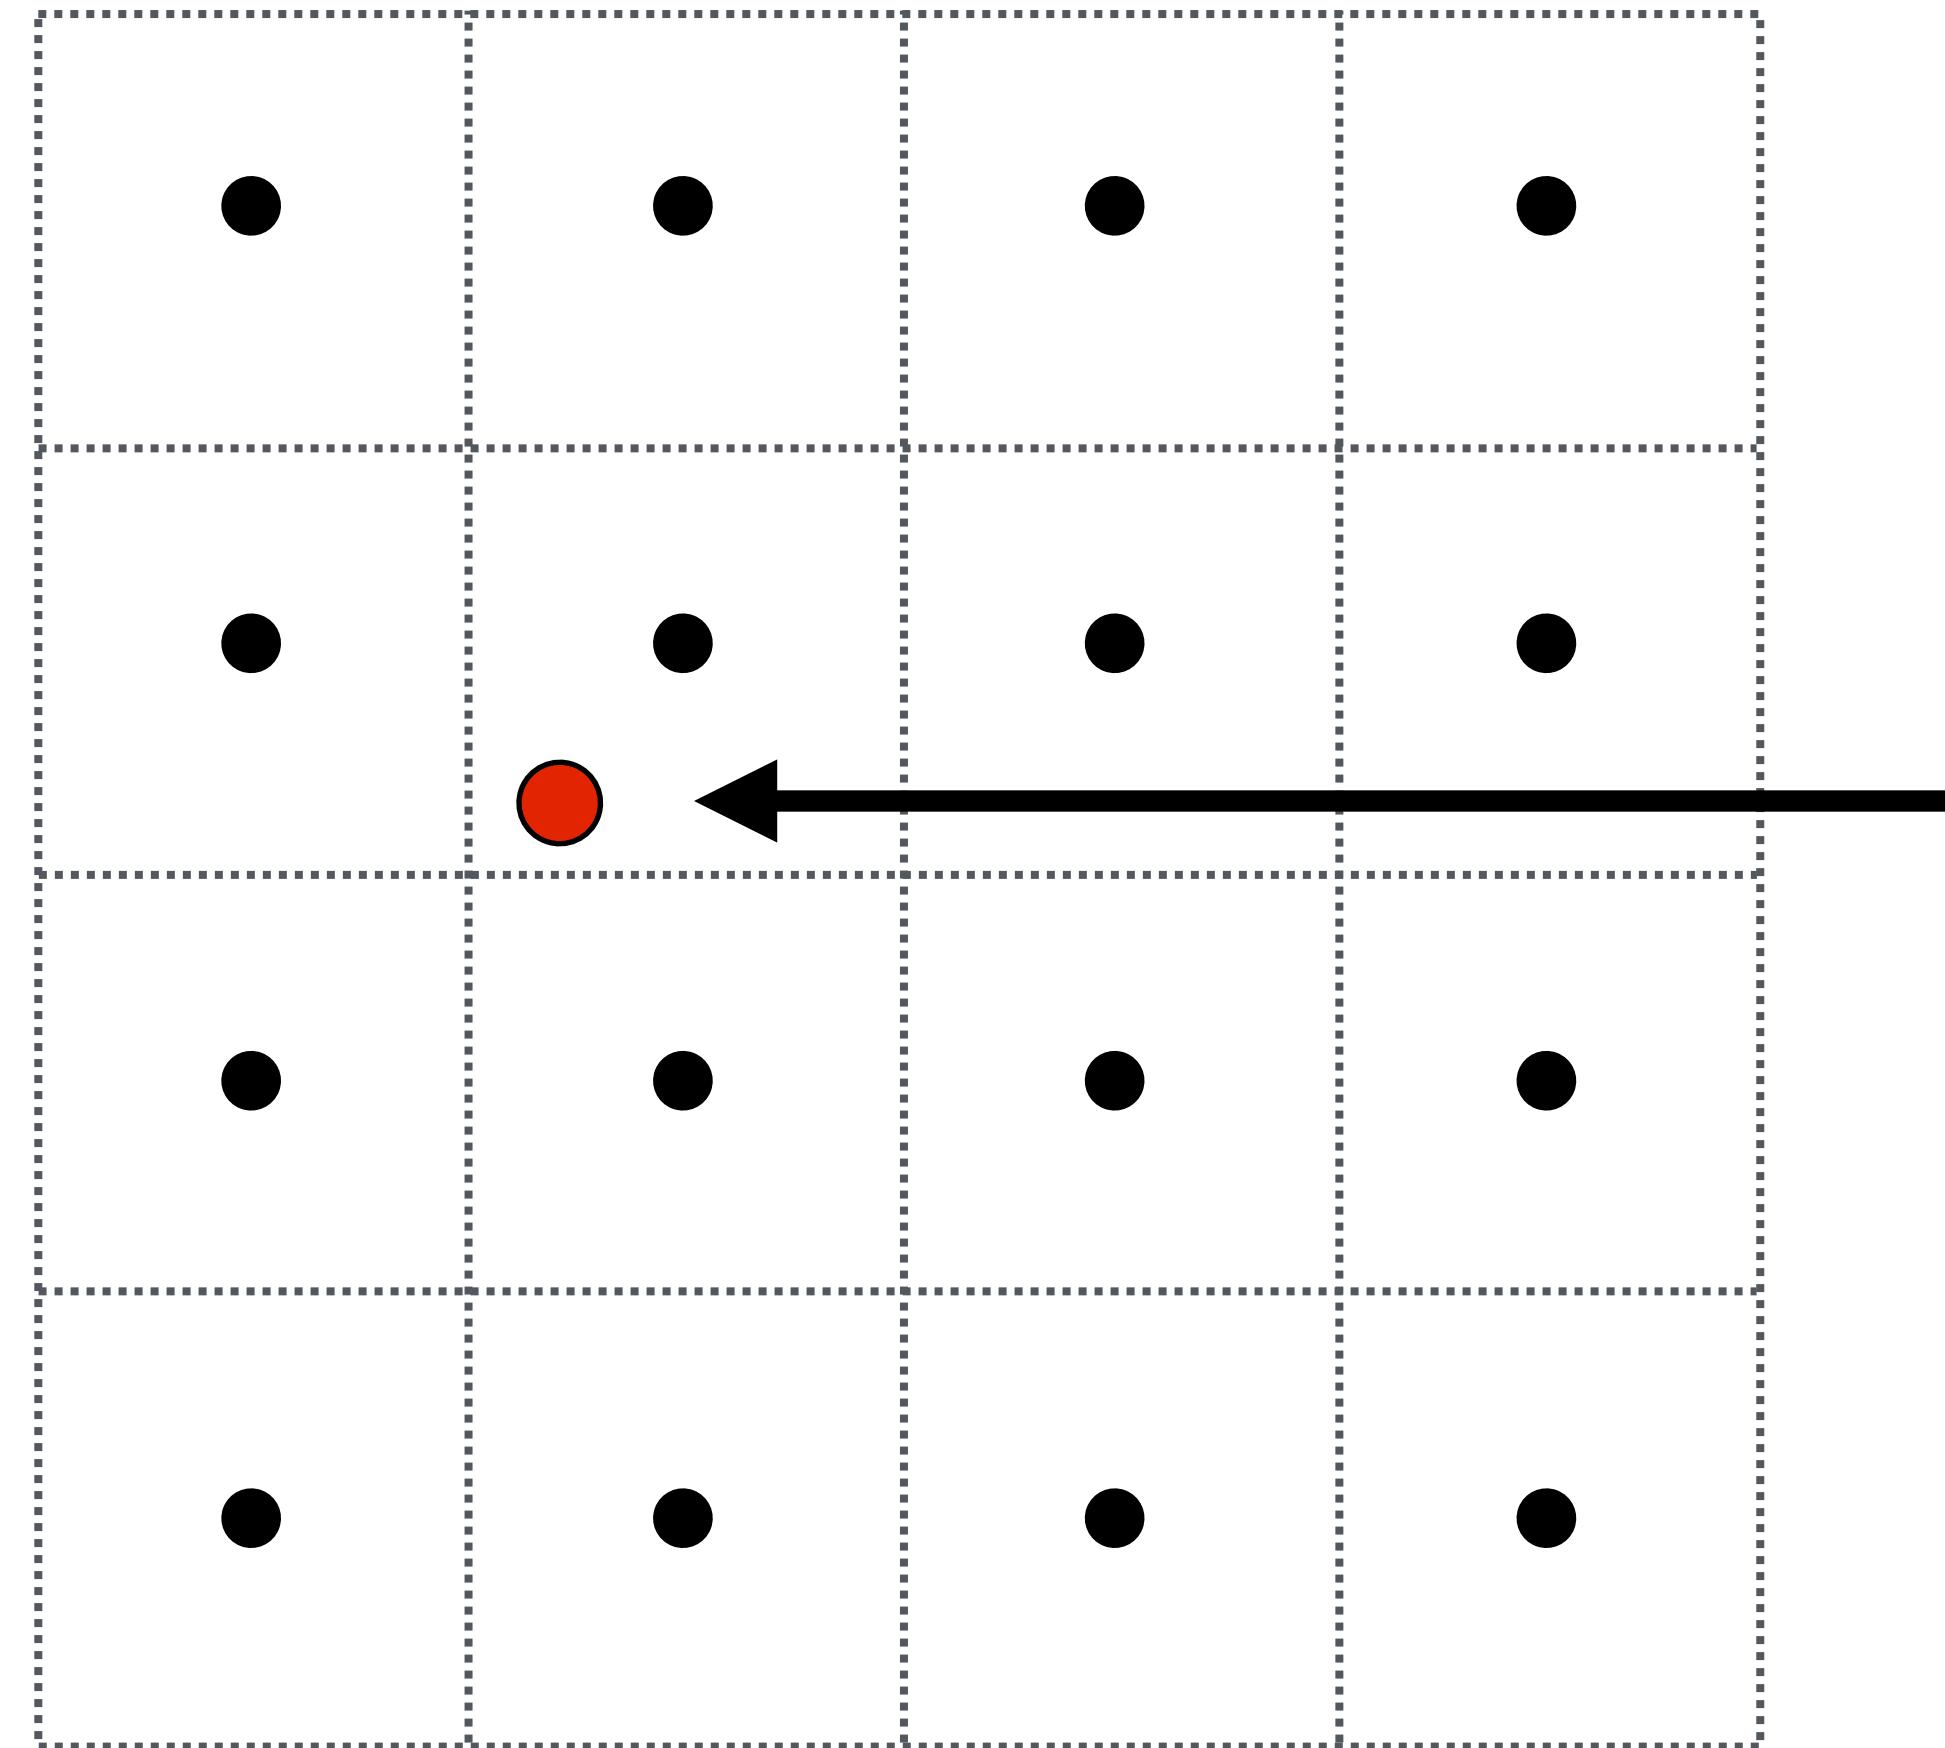
Texture magnification (bilinear)



Texture magnification (bilinear)



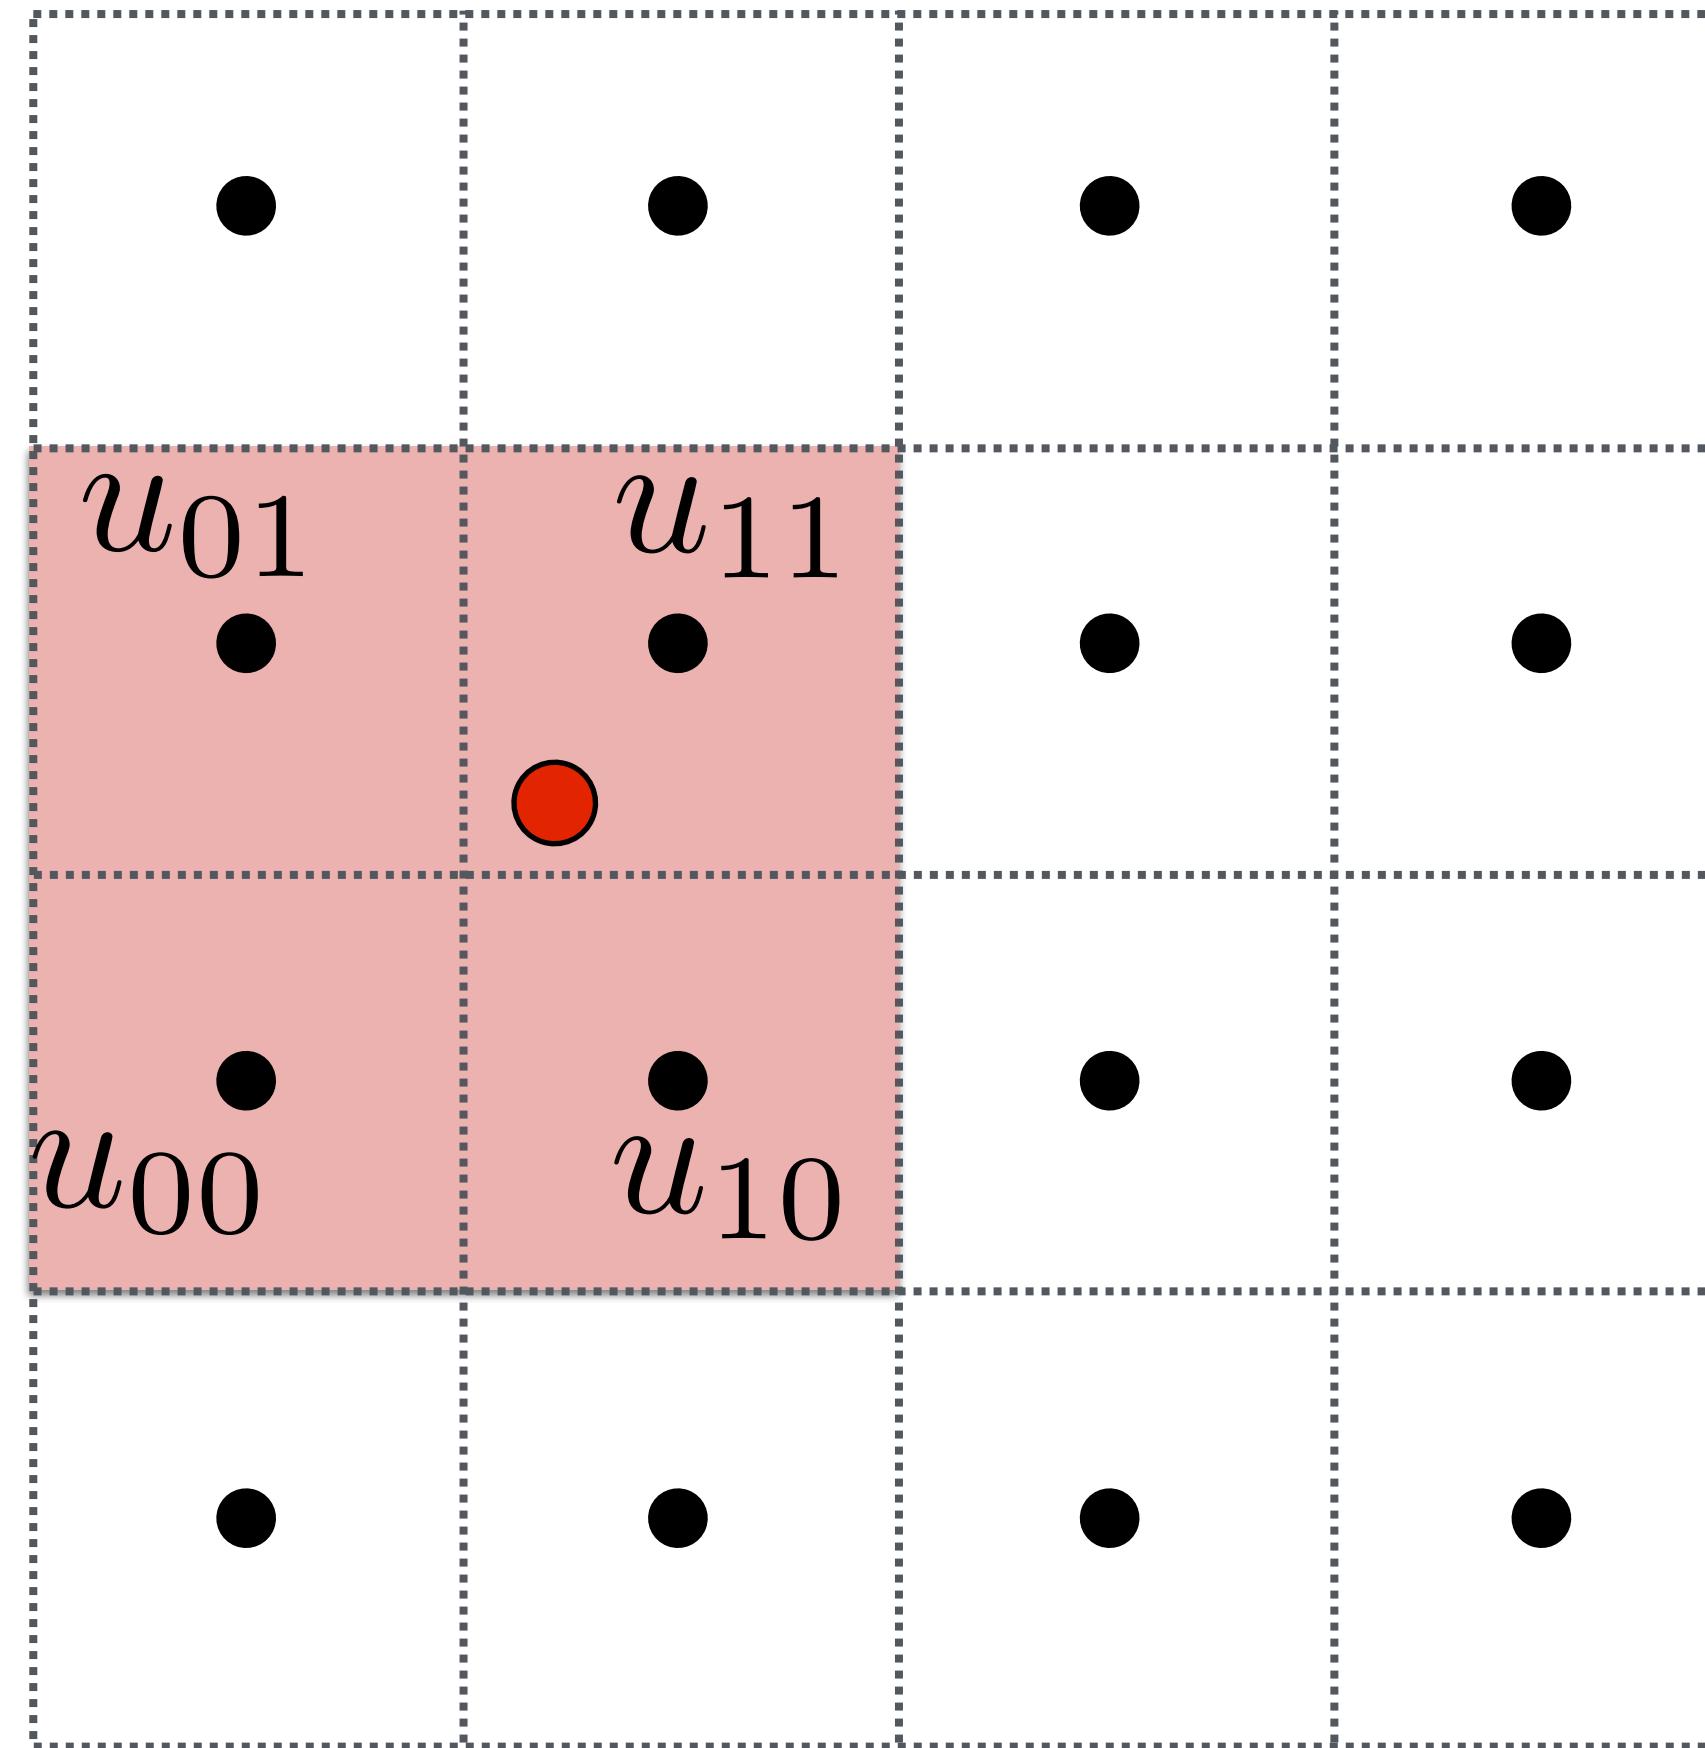
Bilinear interpolation



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

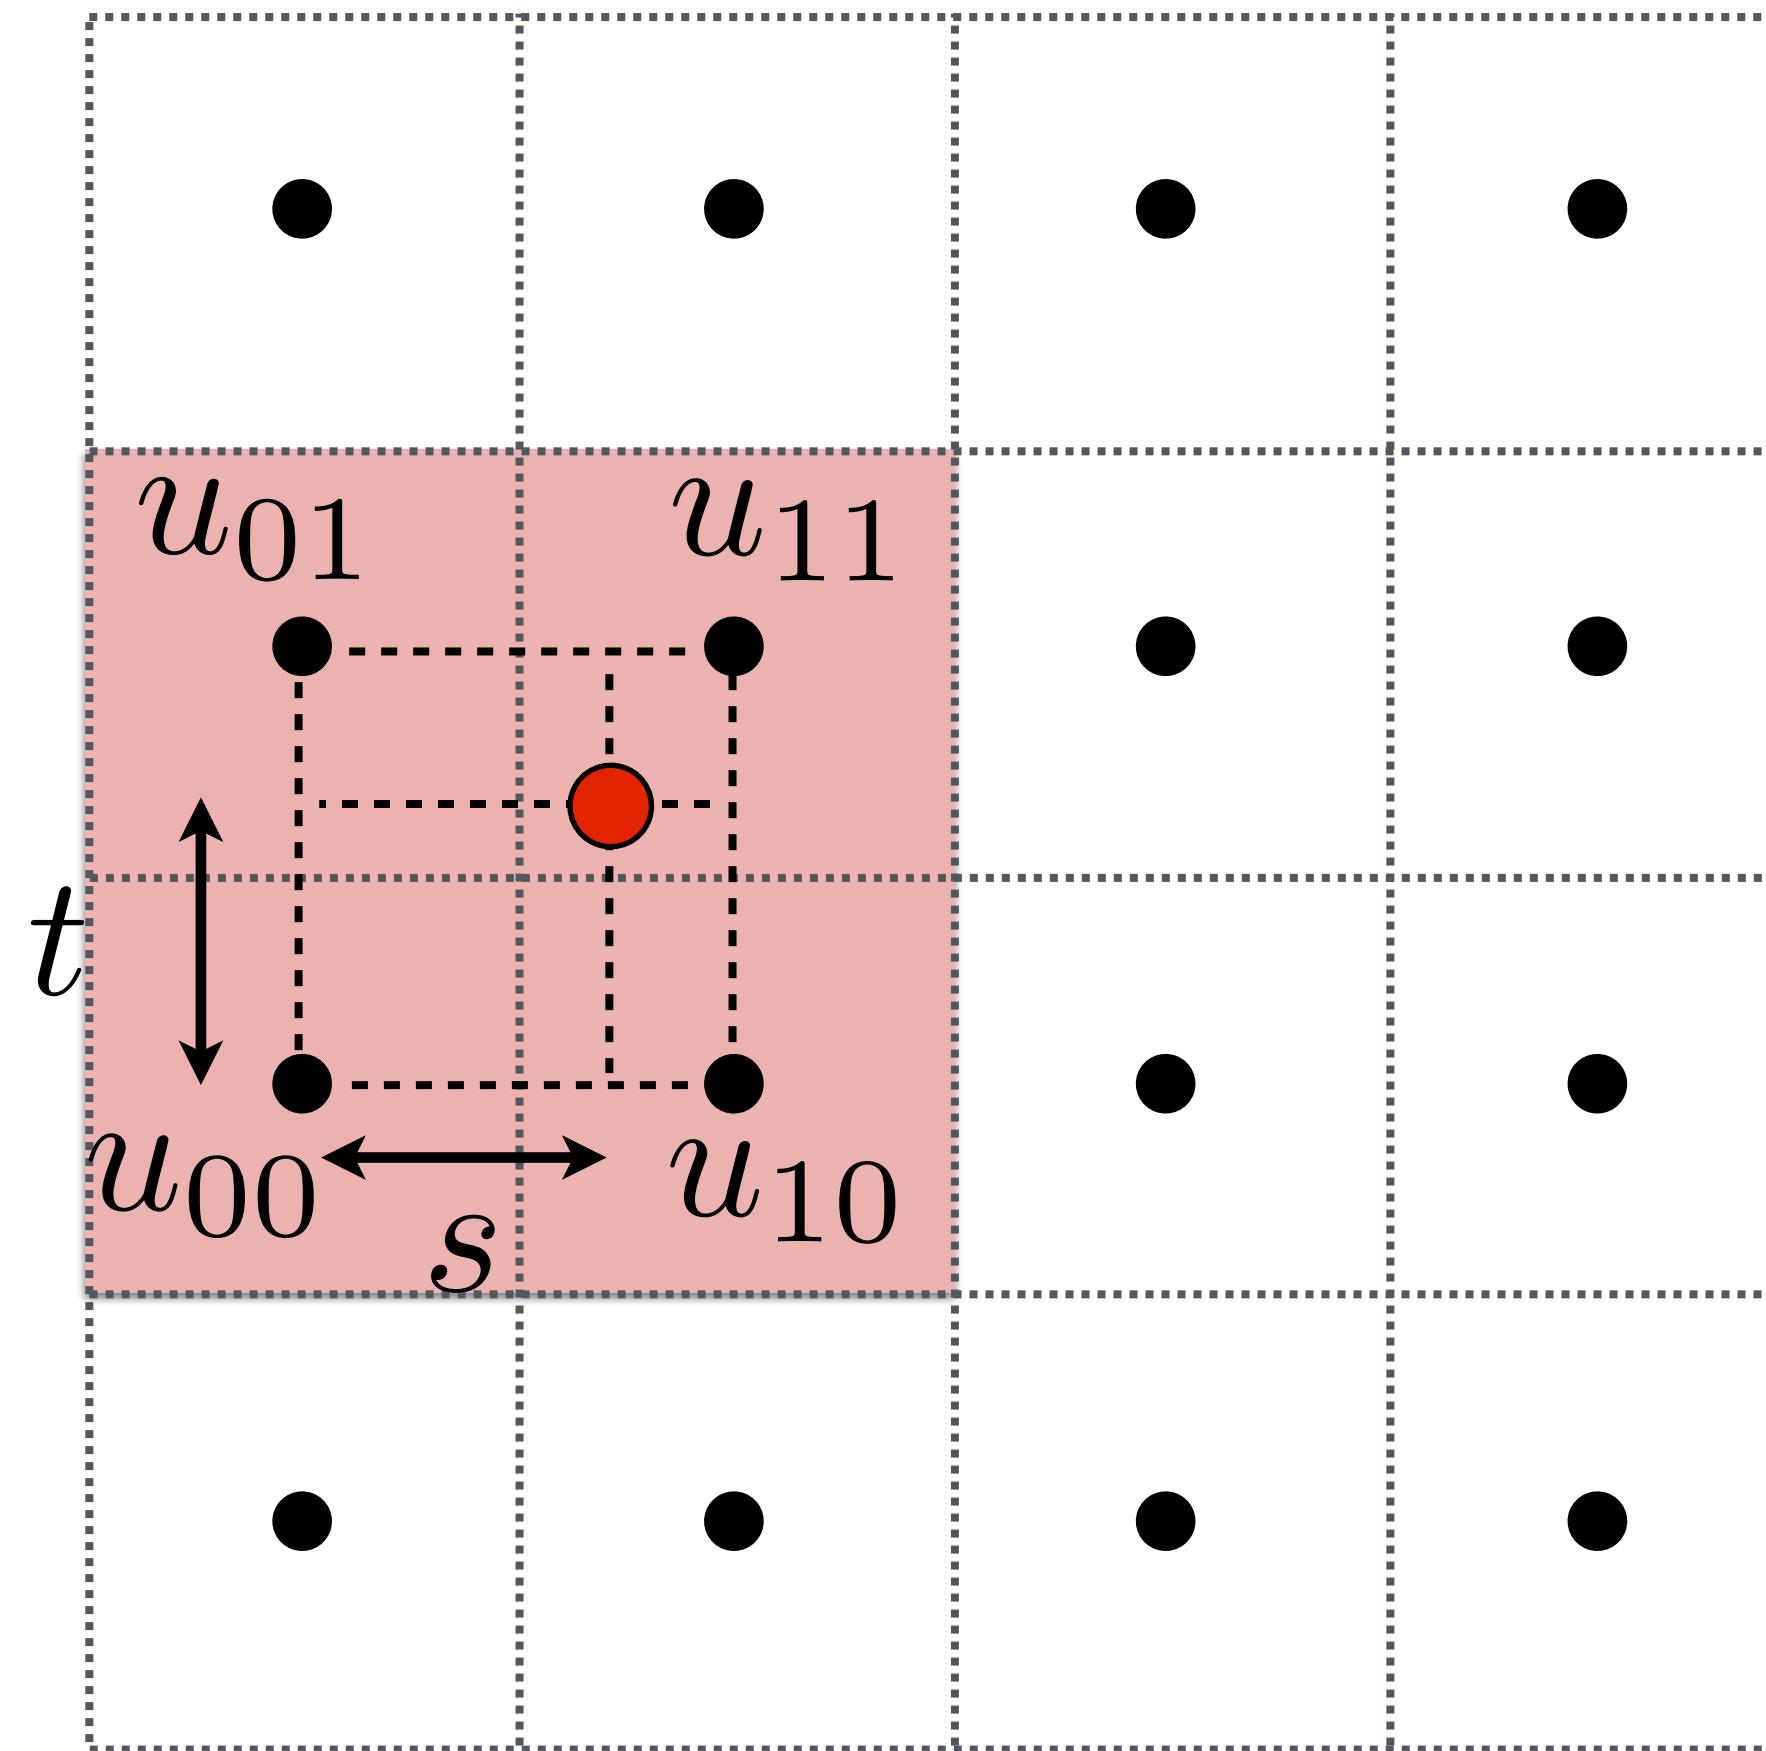
Black points indicate texture sample locations

Bilinear interpolation



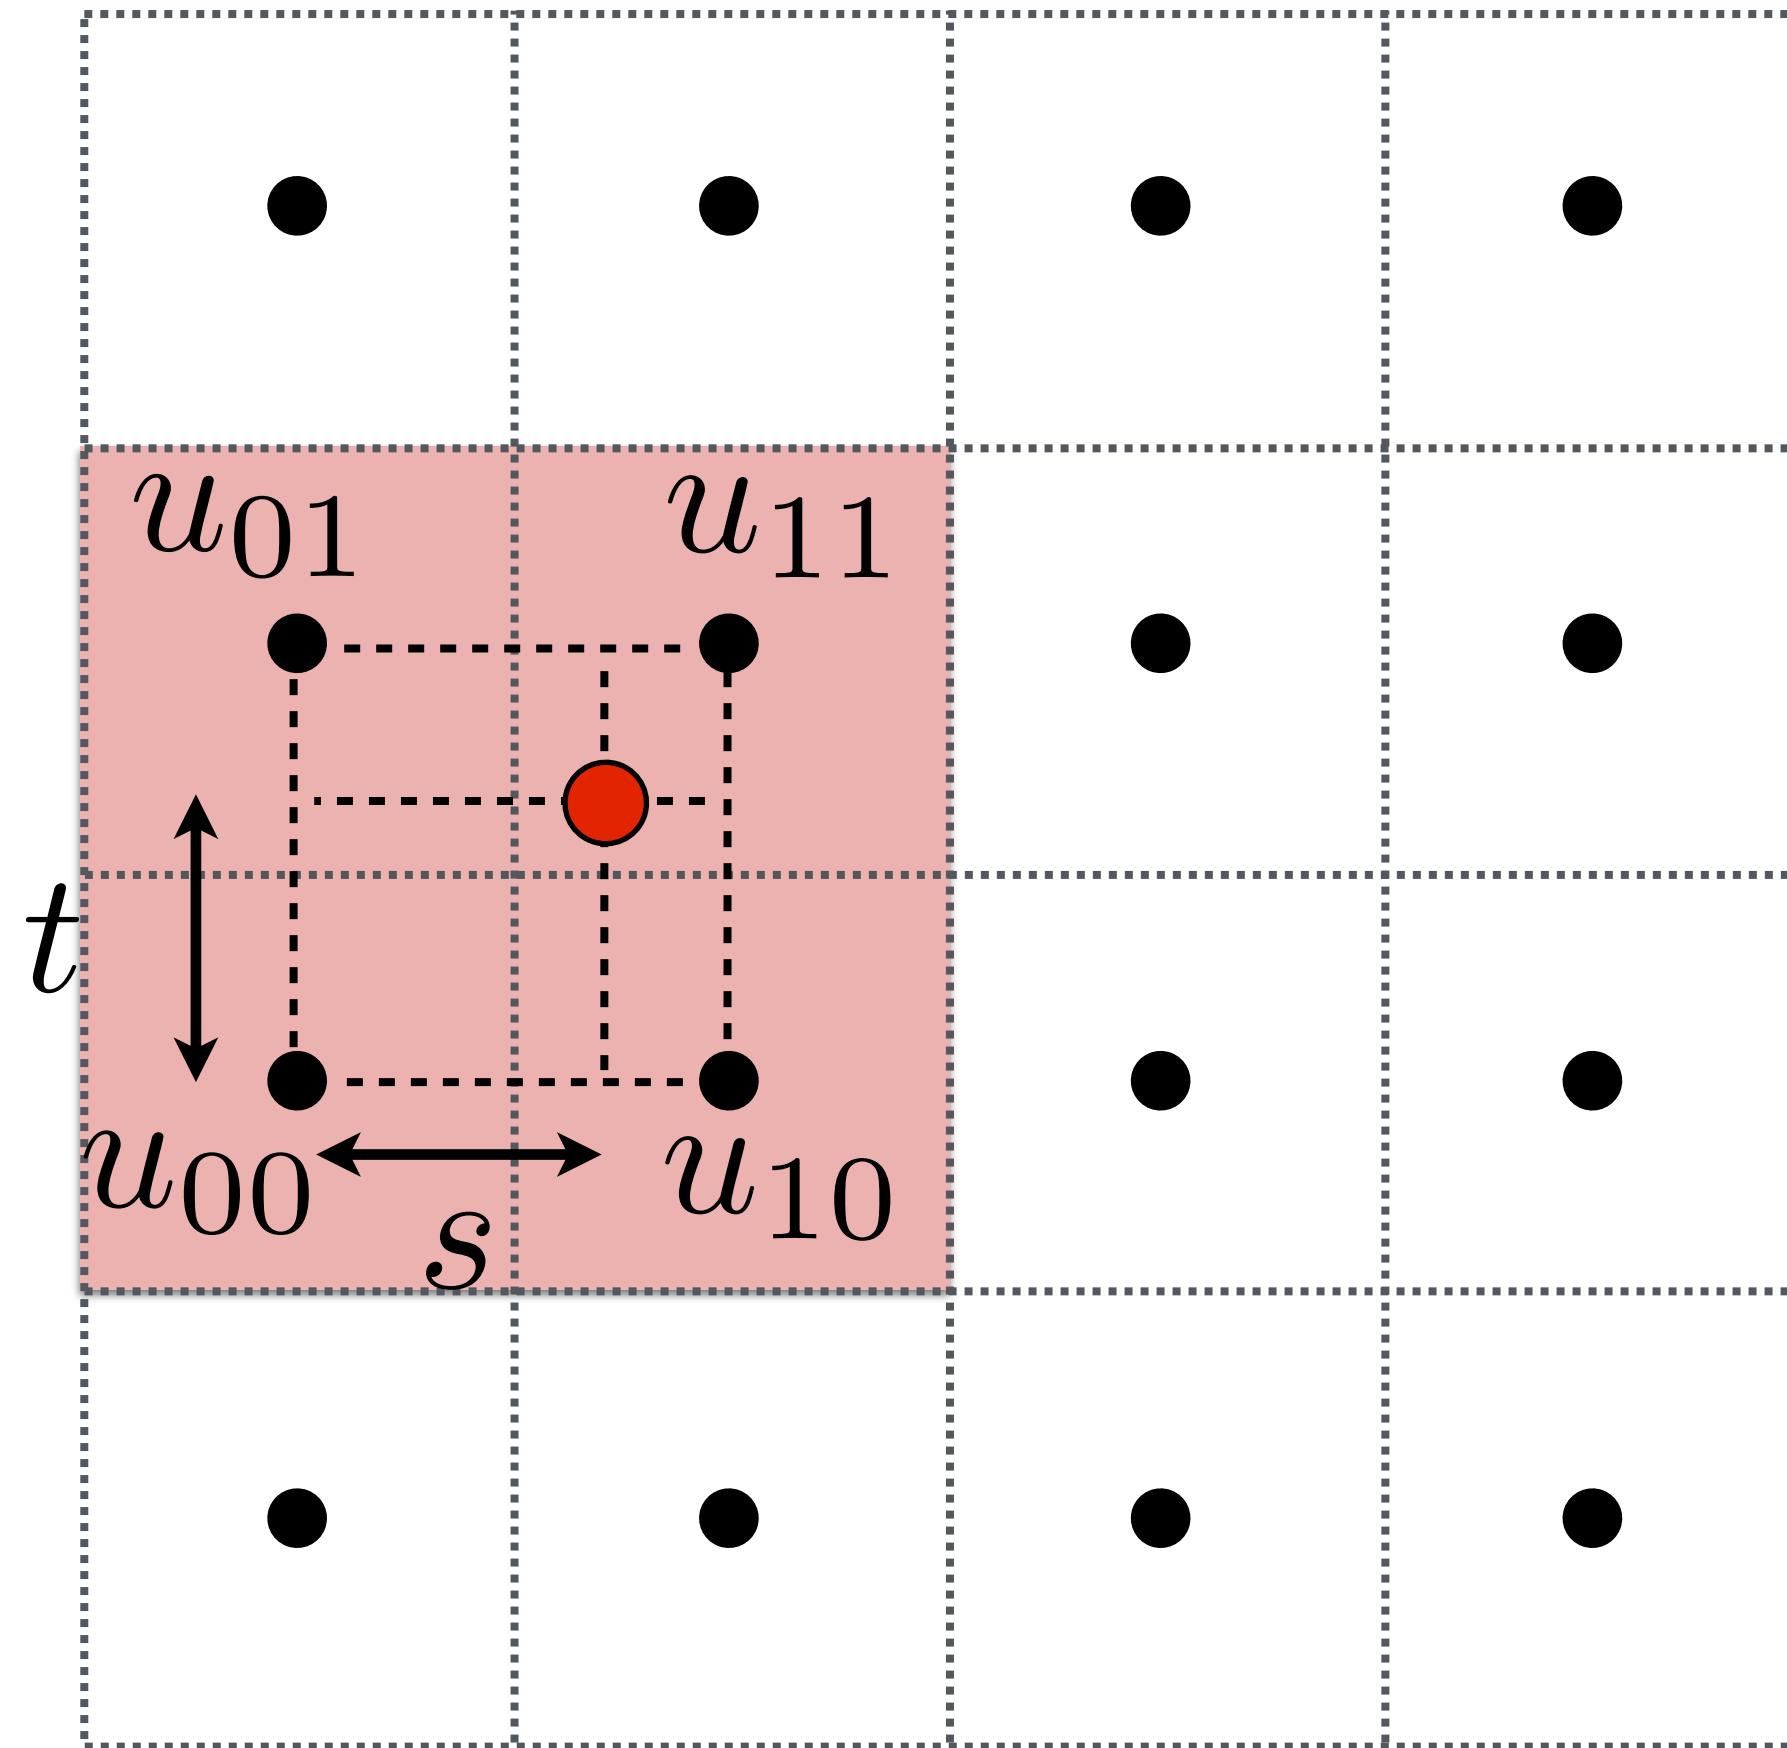
Take 4 nearest sample locations, with texture values as labeled.

Bilinear interpolation



**And fractional offsets,
(s,t) as shown**

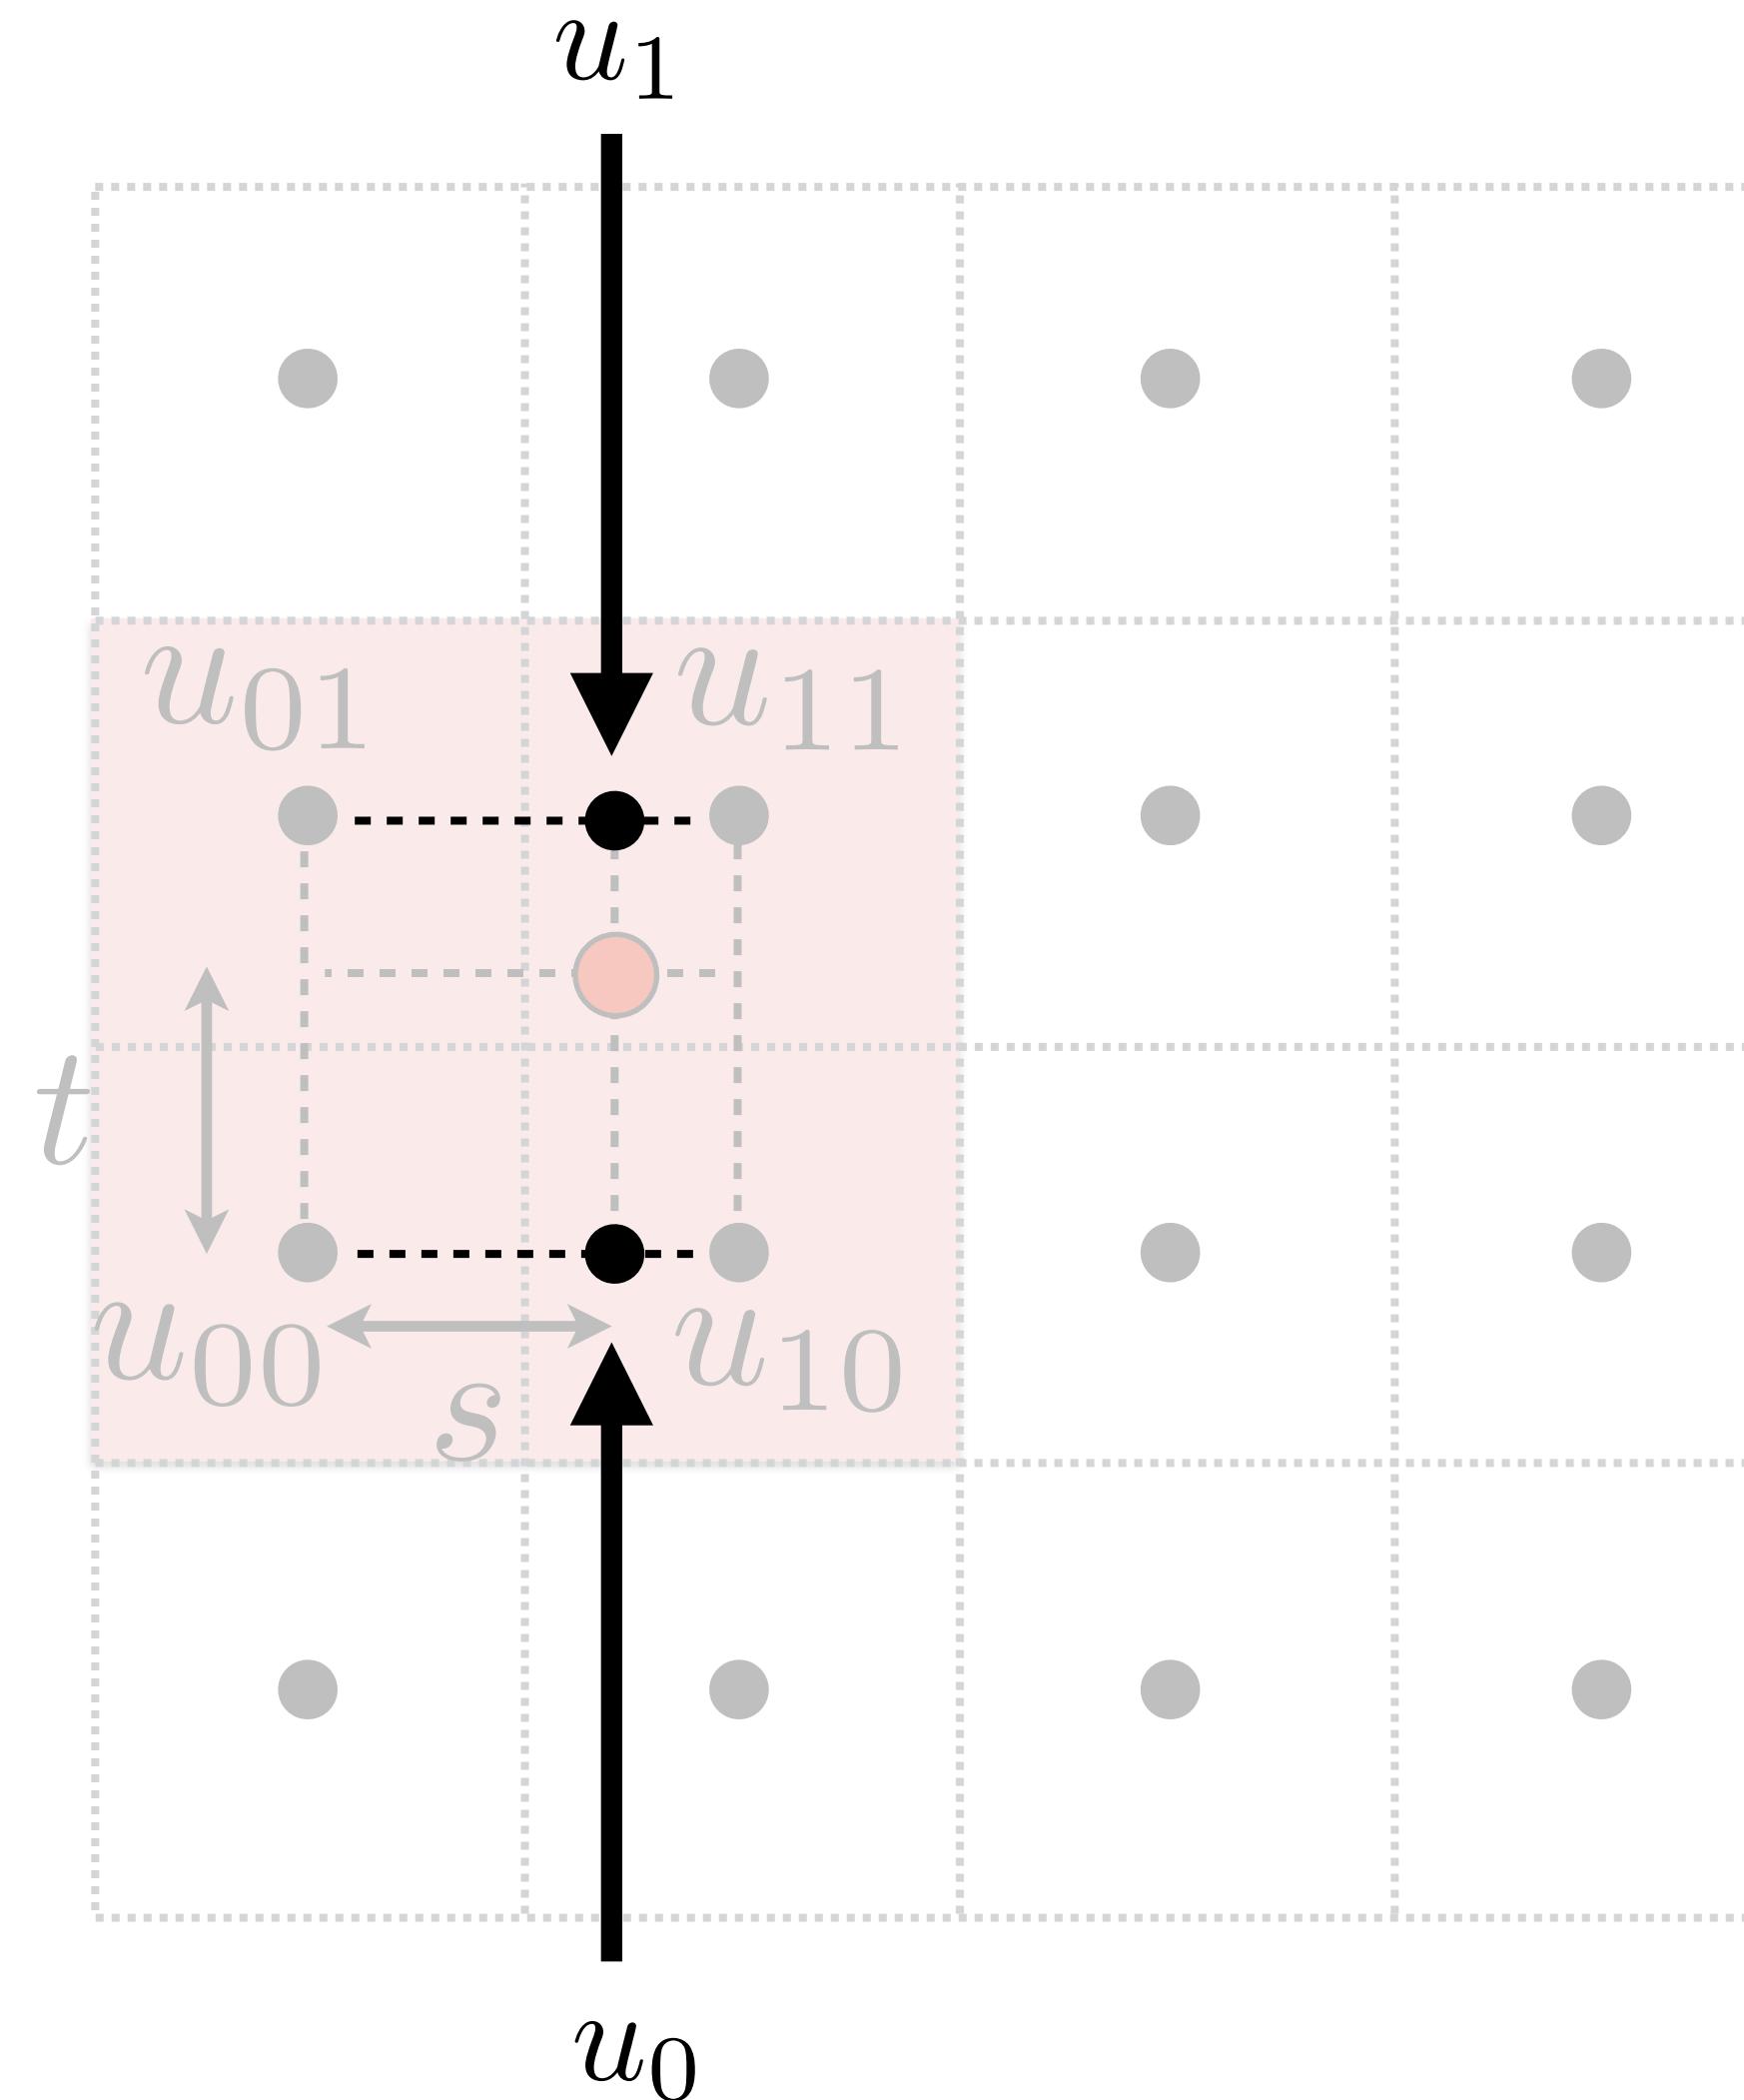
Bilinear interpolation



Linear interpolation (1D)

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

Bilinear interpolation



Linear interpolation (1D)

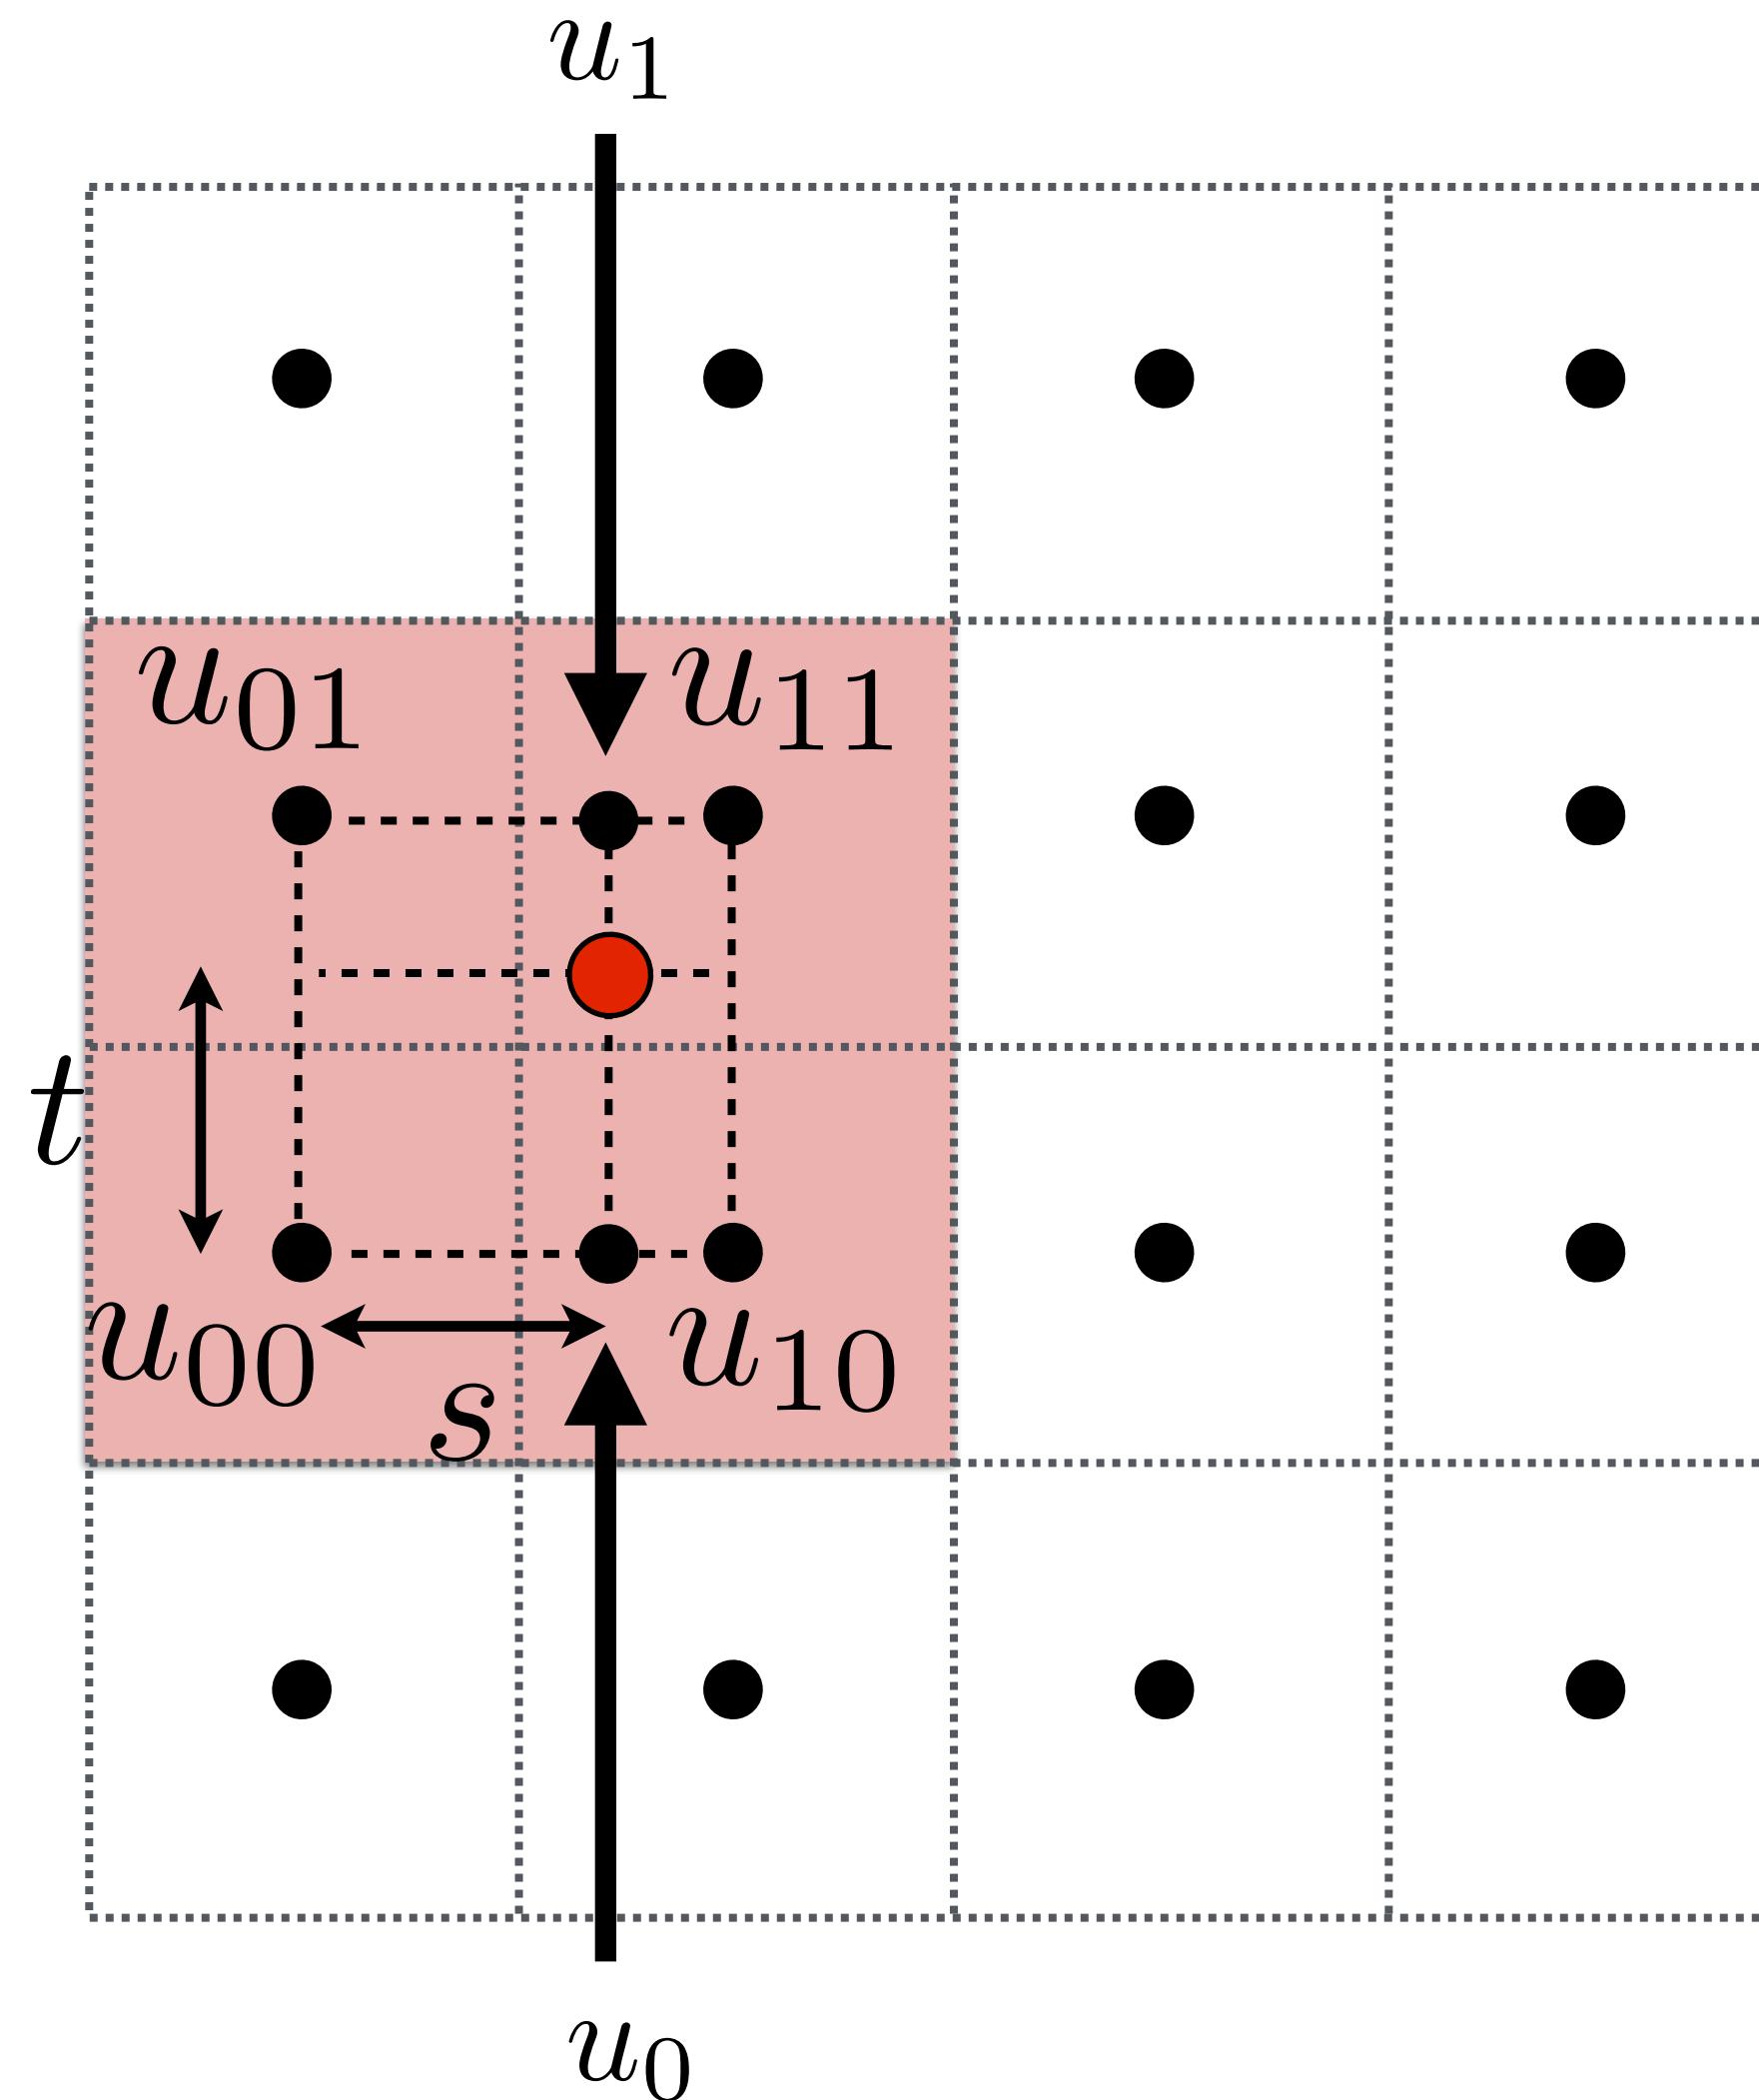
$$\text{lerp}(x, v_0, v_1) = v_0 + x(v_1 - v_0)$$

Two helper lerps (horizontal)

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

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

Bilinear interpolation



Linear interpolation (1D)

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

Two helper lerps

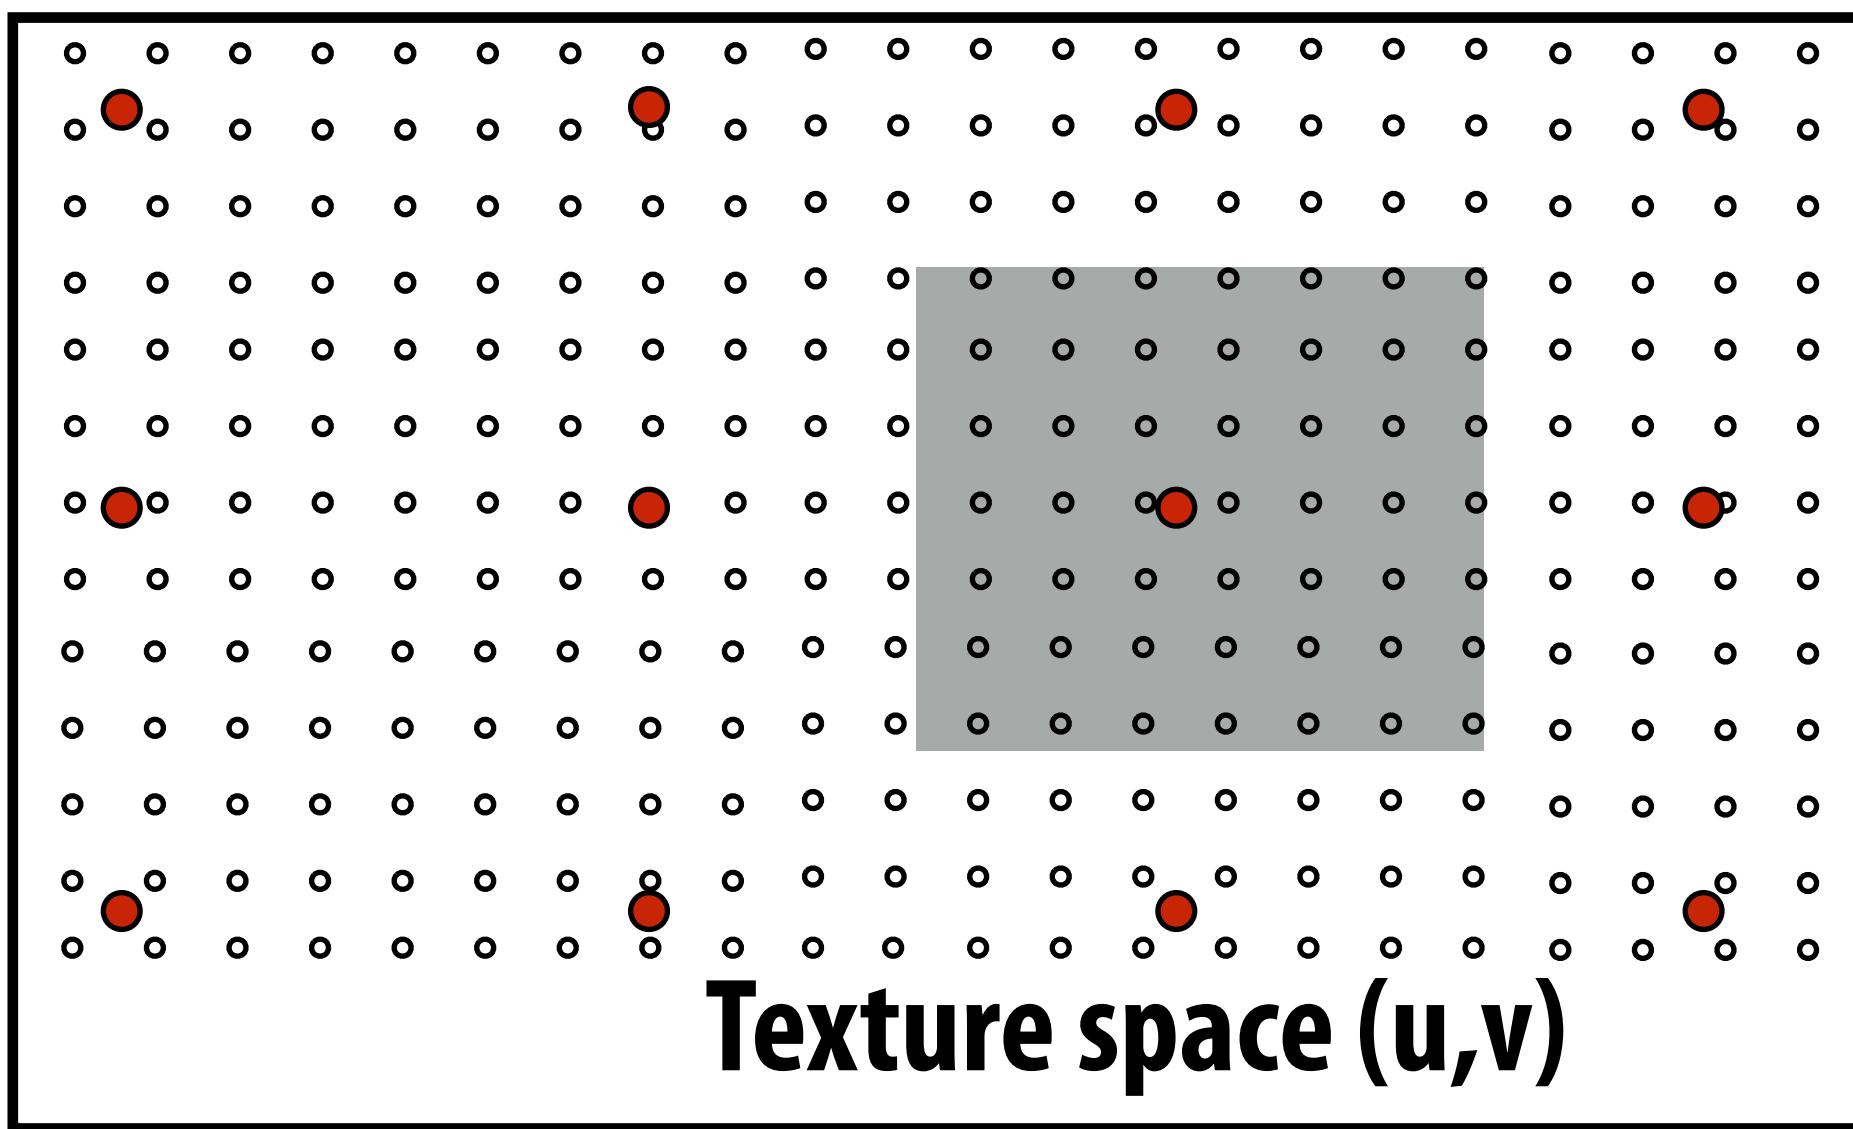
$$u_0 = \text{lerp}(s, u_{00}, u_{10})$$

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

Final vertical lerp, to get result:

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

Texture minification



By now I hope you've realized:

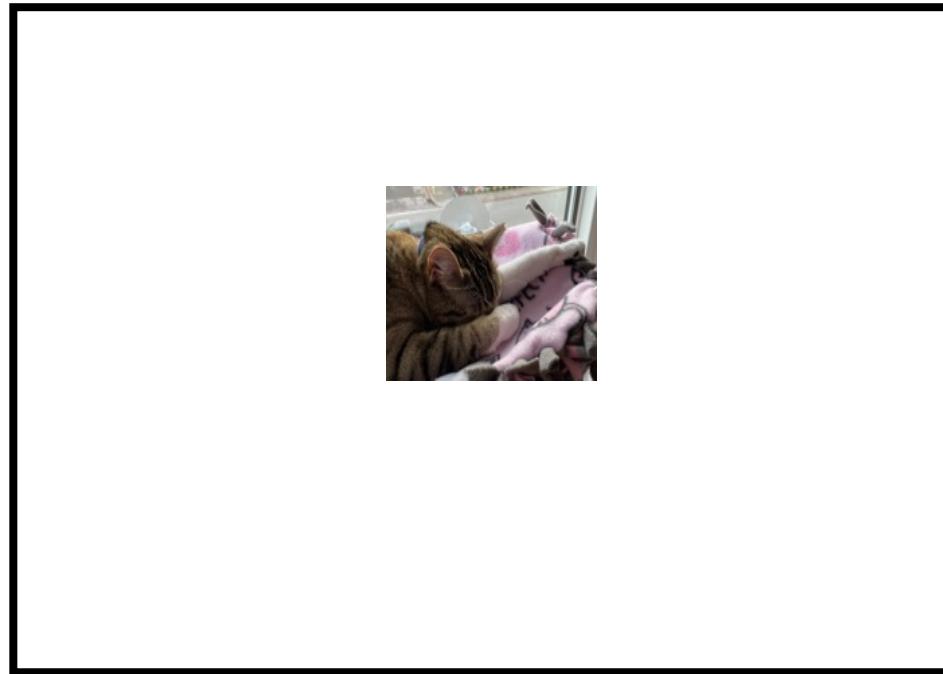
Applying textures is a form of sampling!

$$t(u, v)$$

Minification of Josephine

Imagine the texture map is 9x9

And is applied to a quad that spans a 3x3 pixel region of screen.



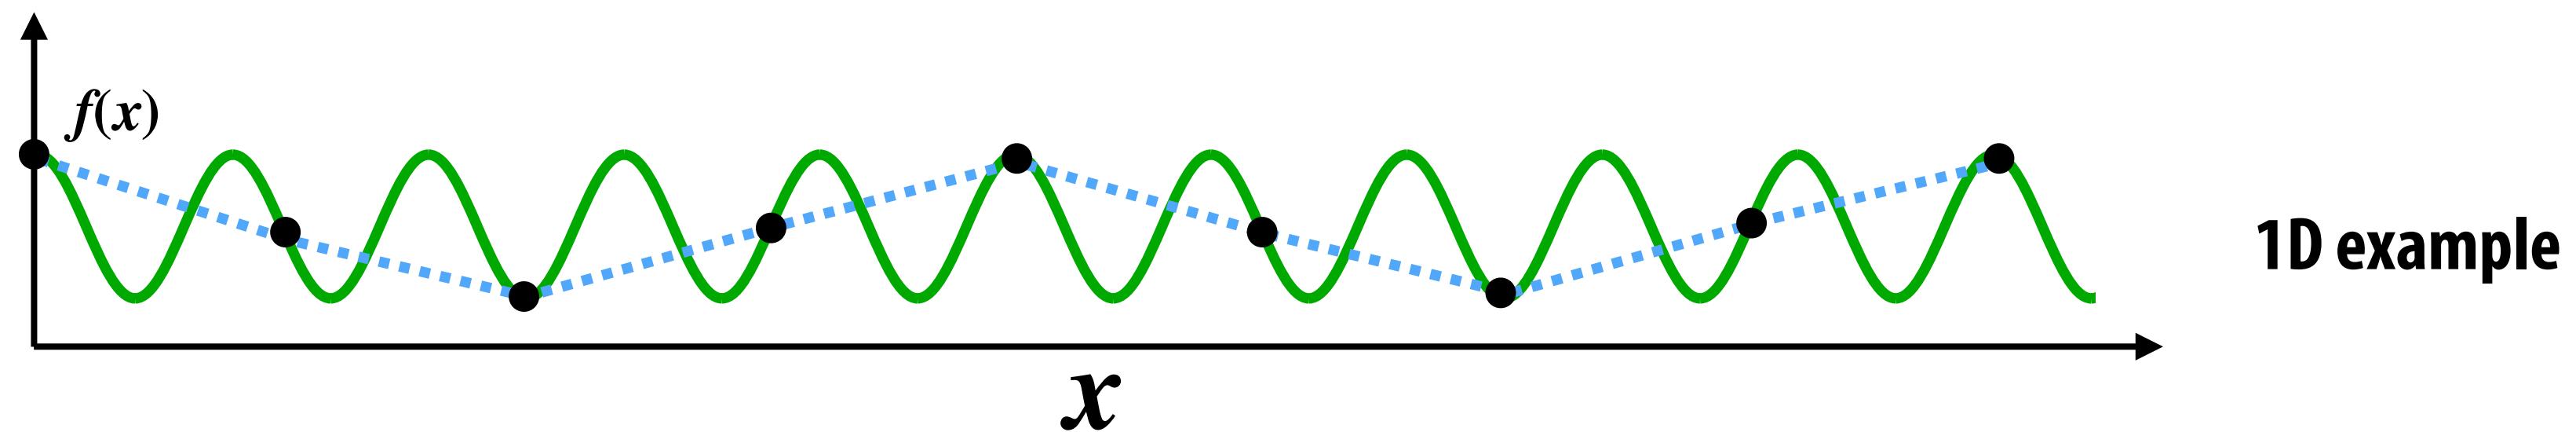
Red dots = samples needed to render

White dots = samples existing in texture map

When a texture is minimized, the texture map is sampled sparsely!

Recall: aliasing

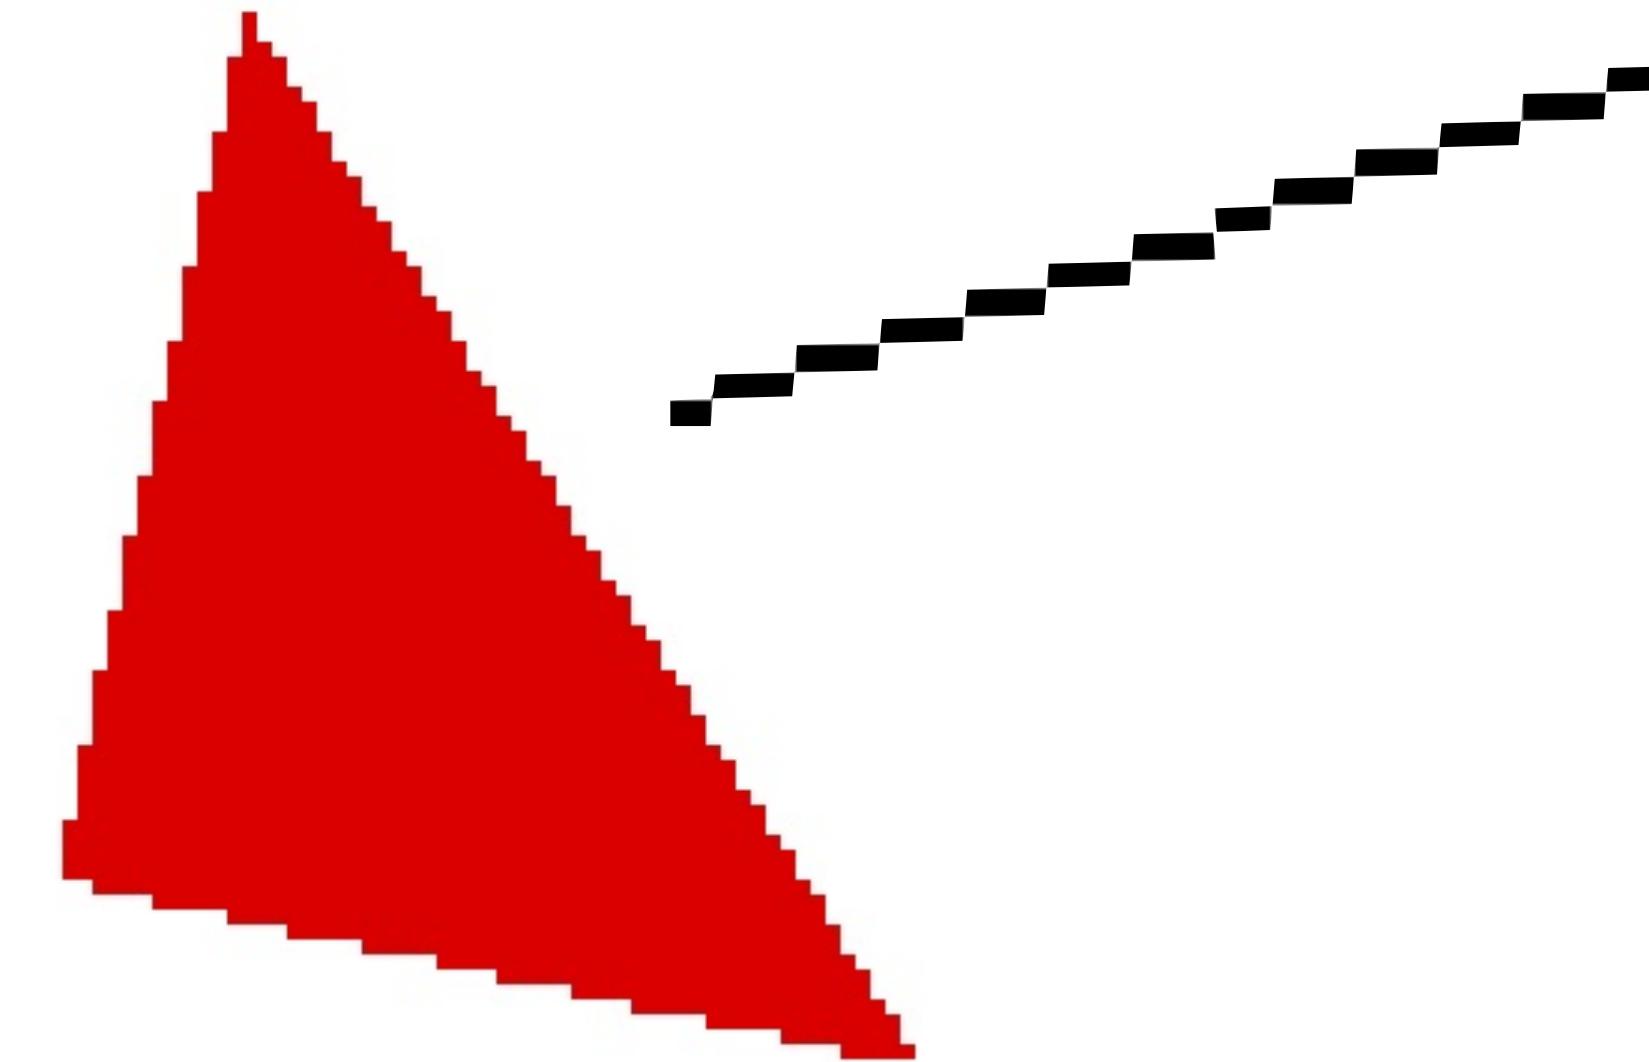
Undersampling a high-frequency signal can result in aliasing



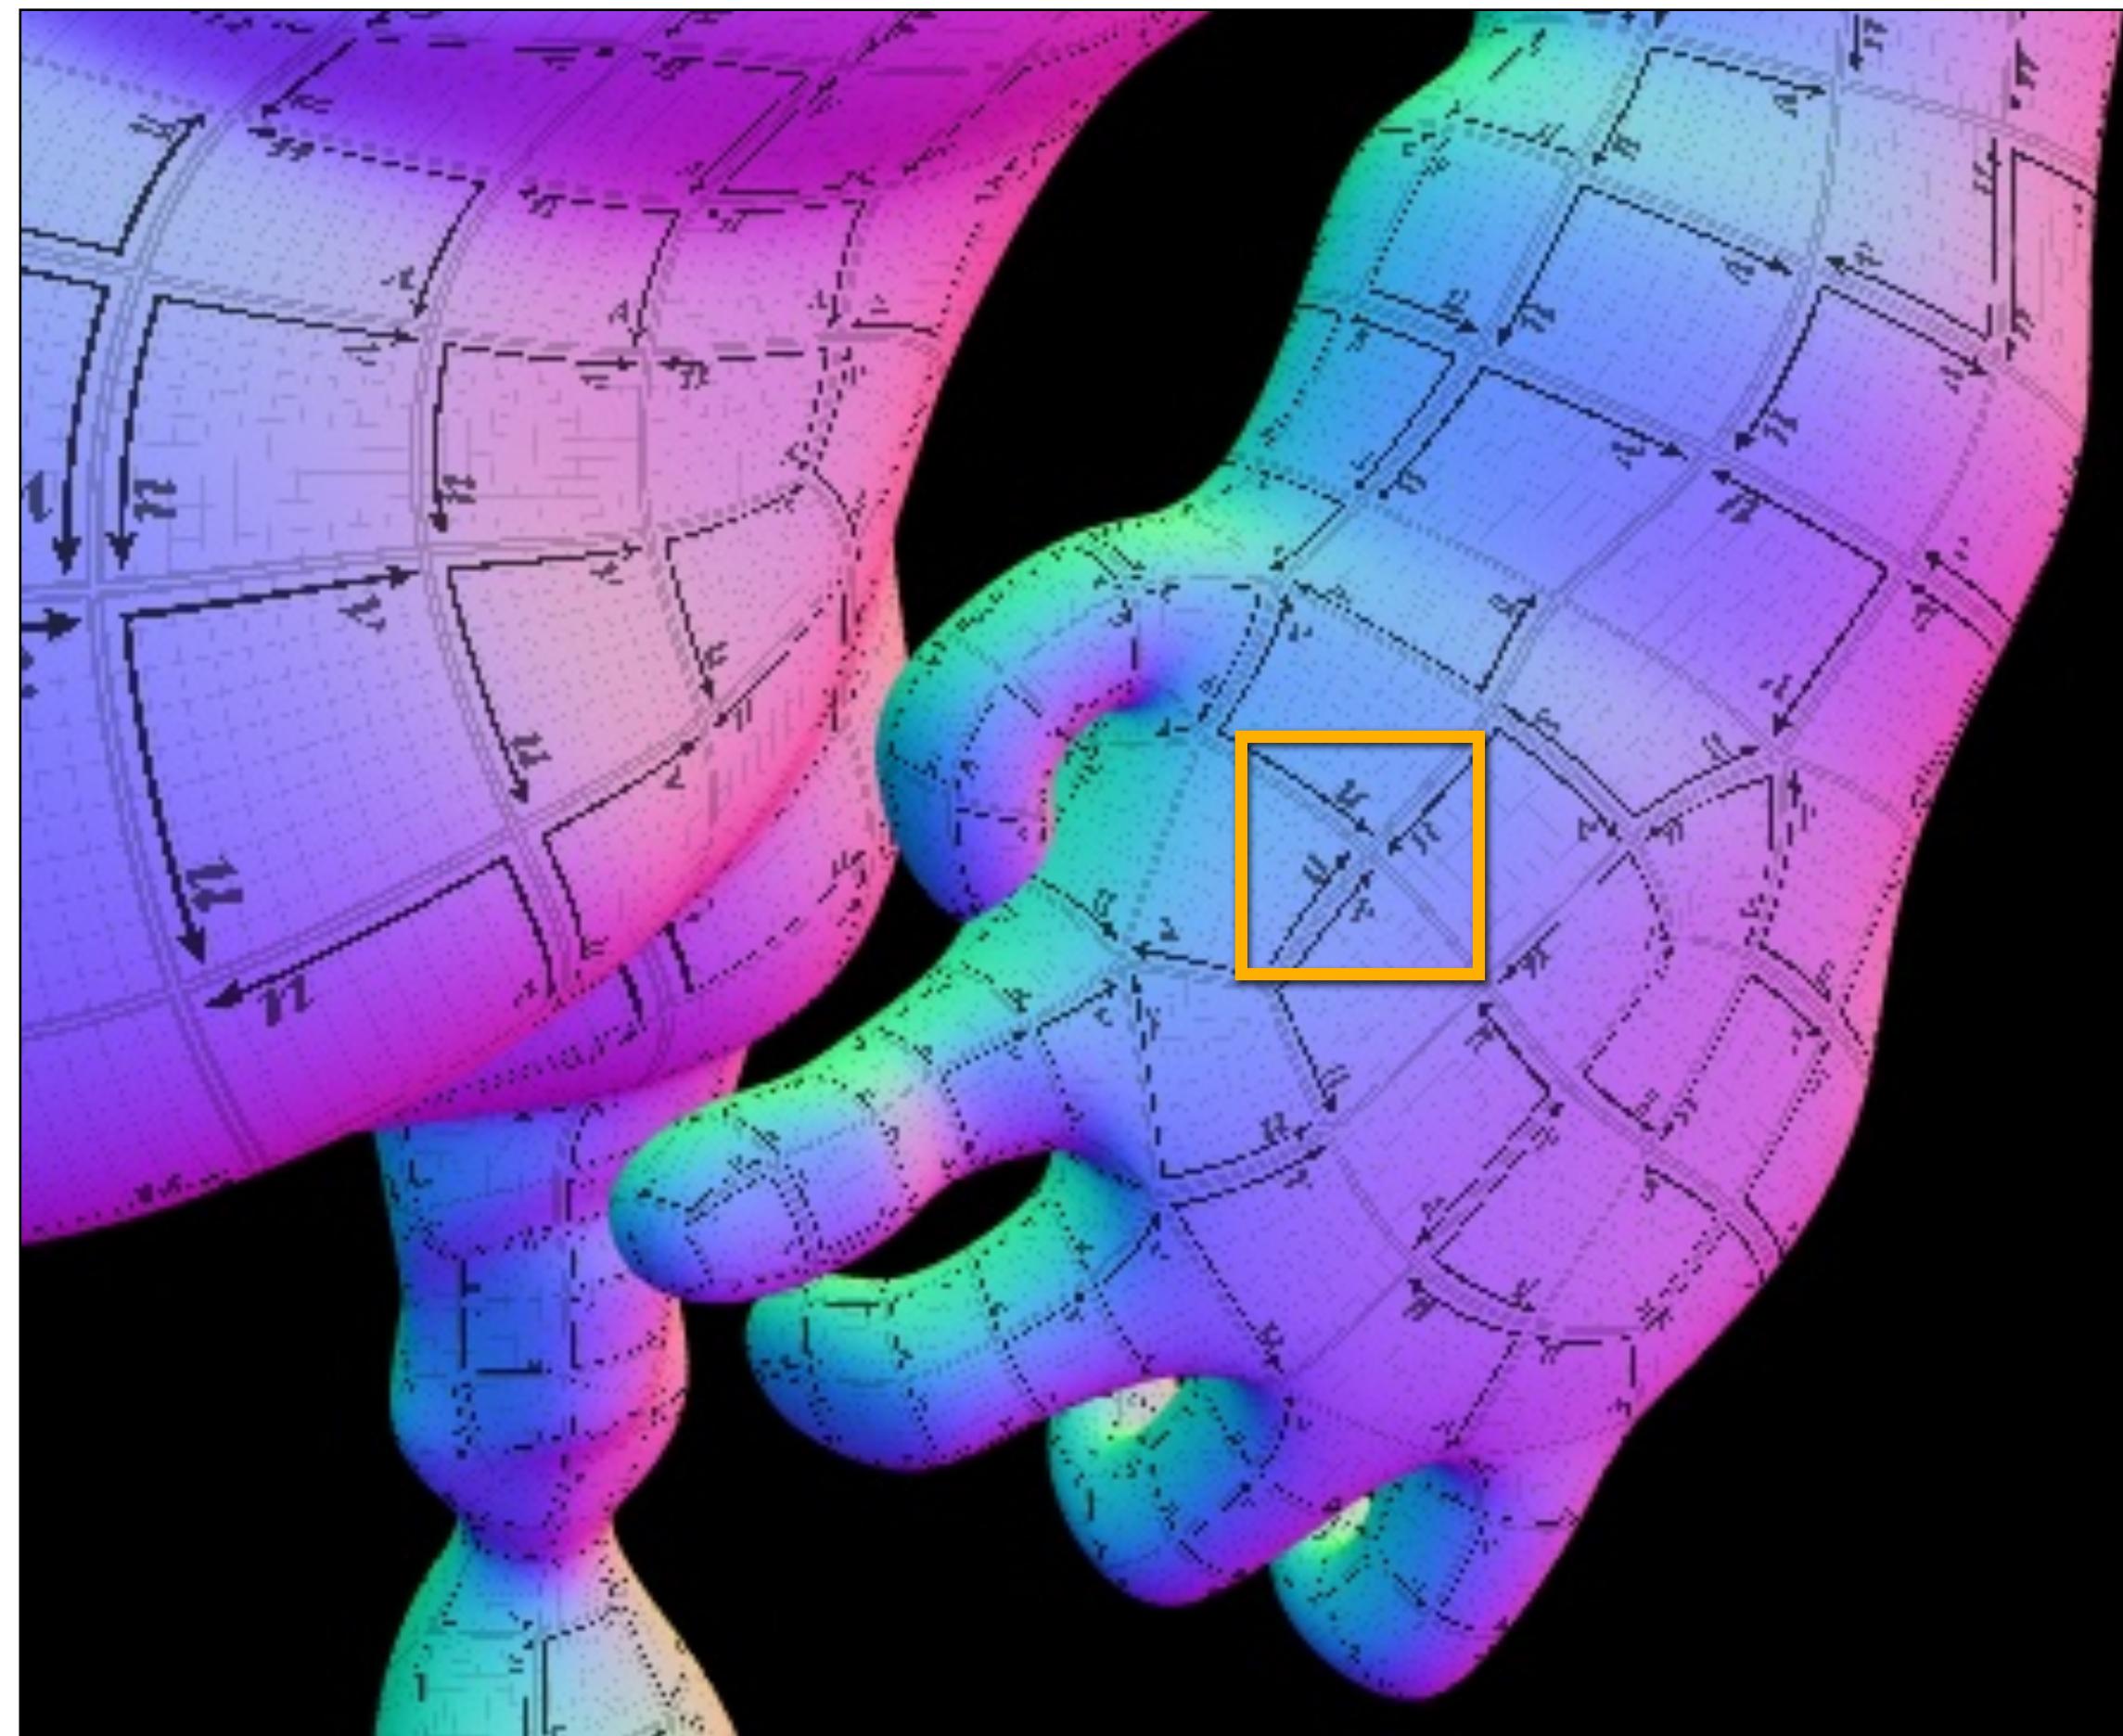
1D example



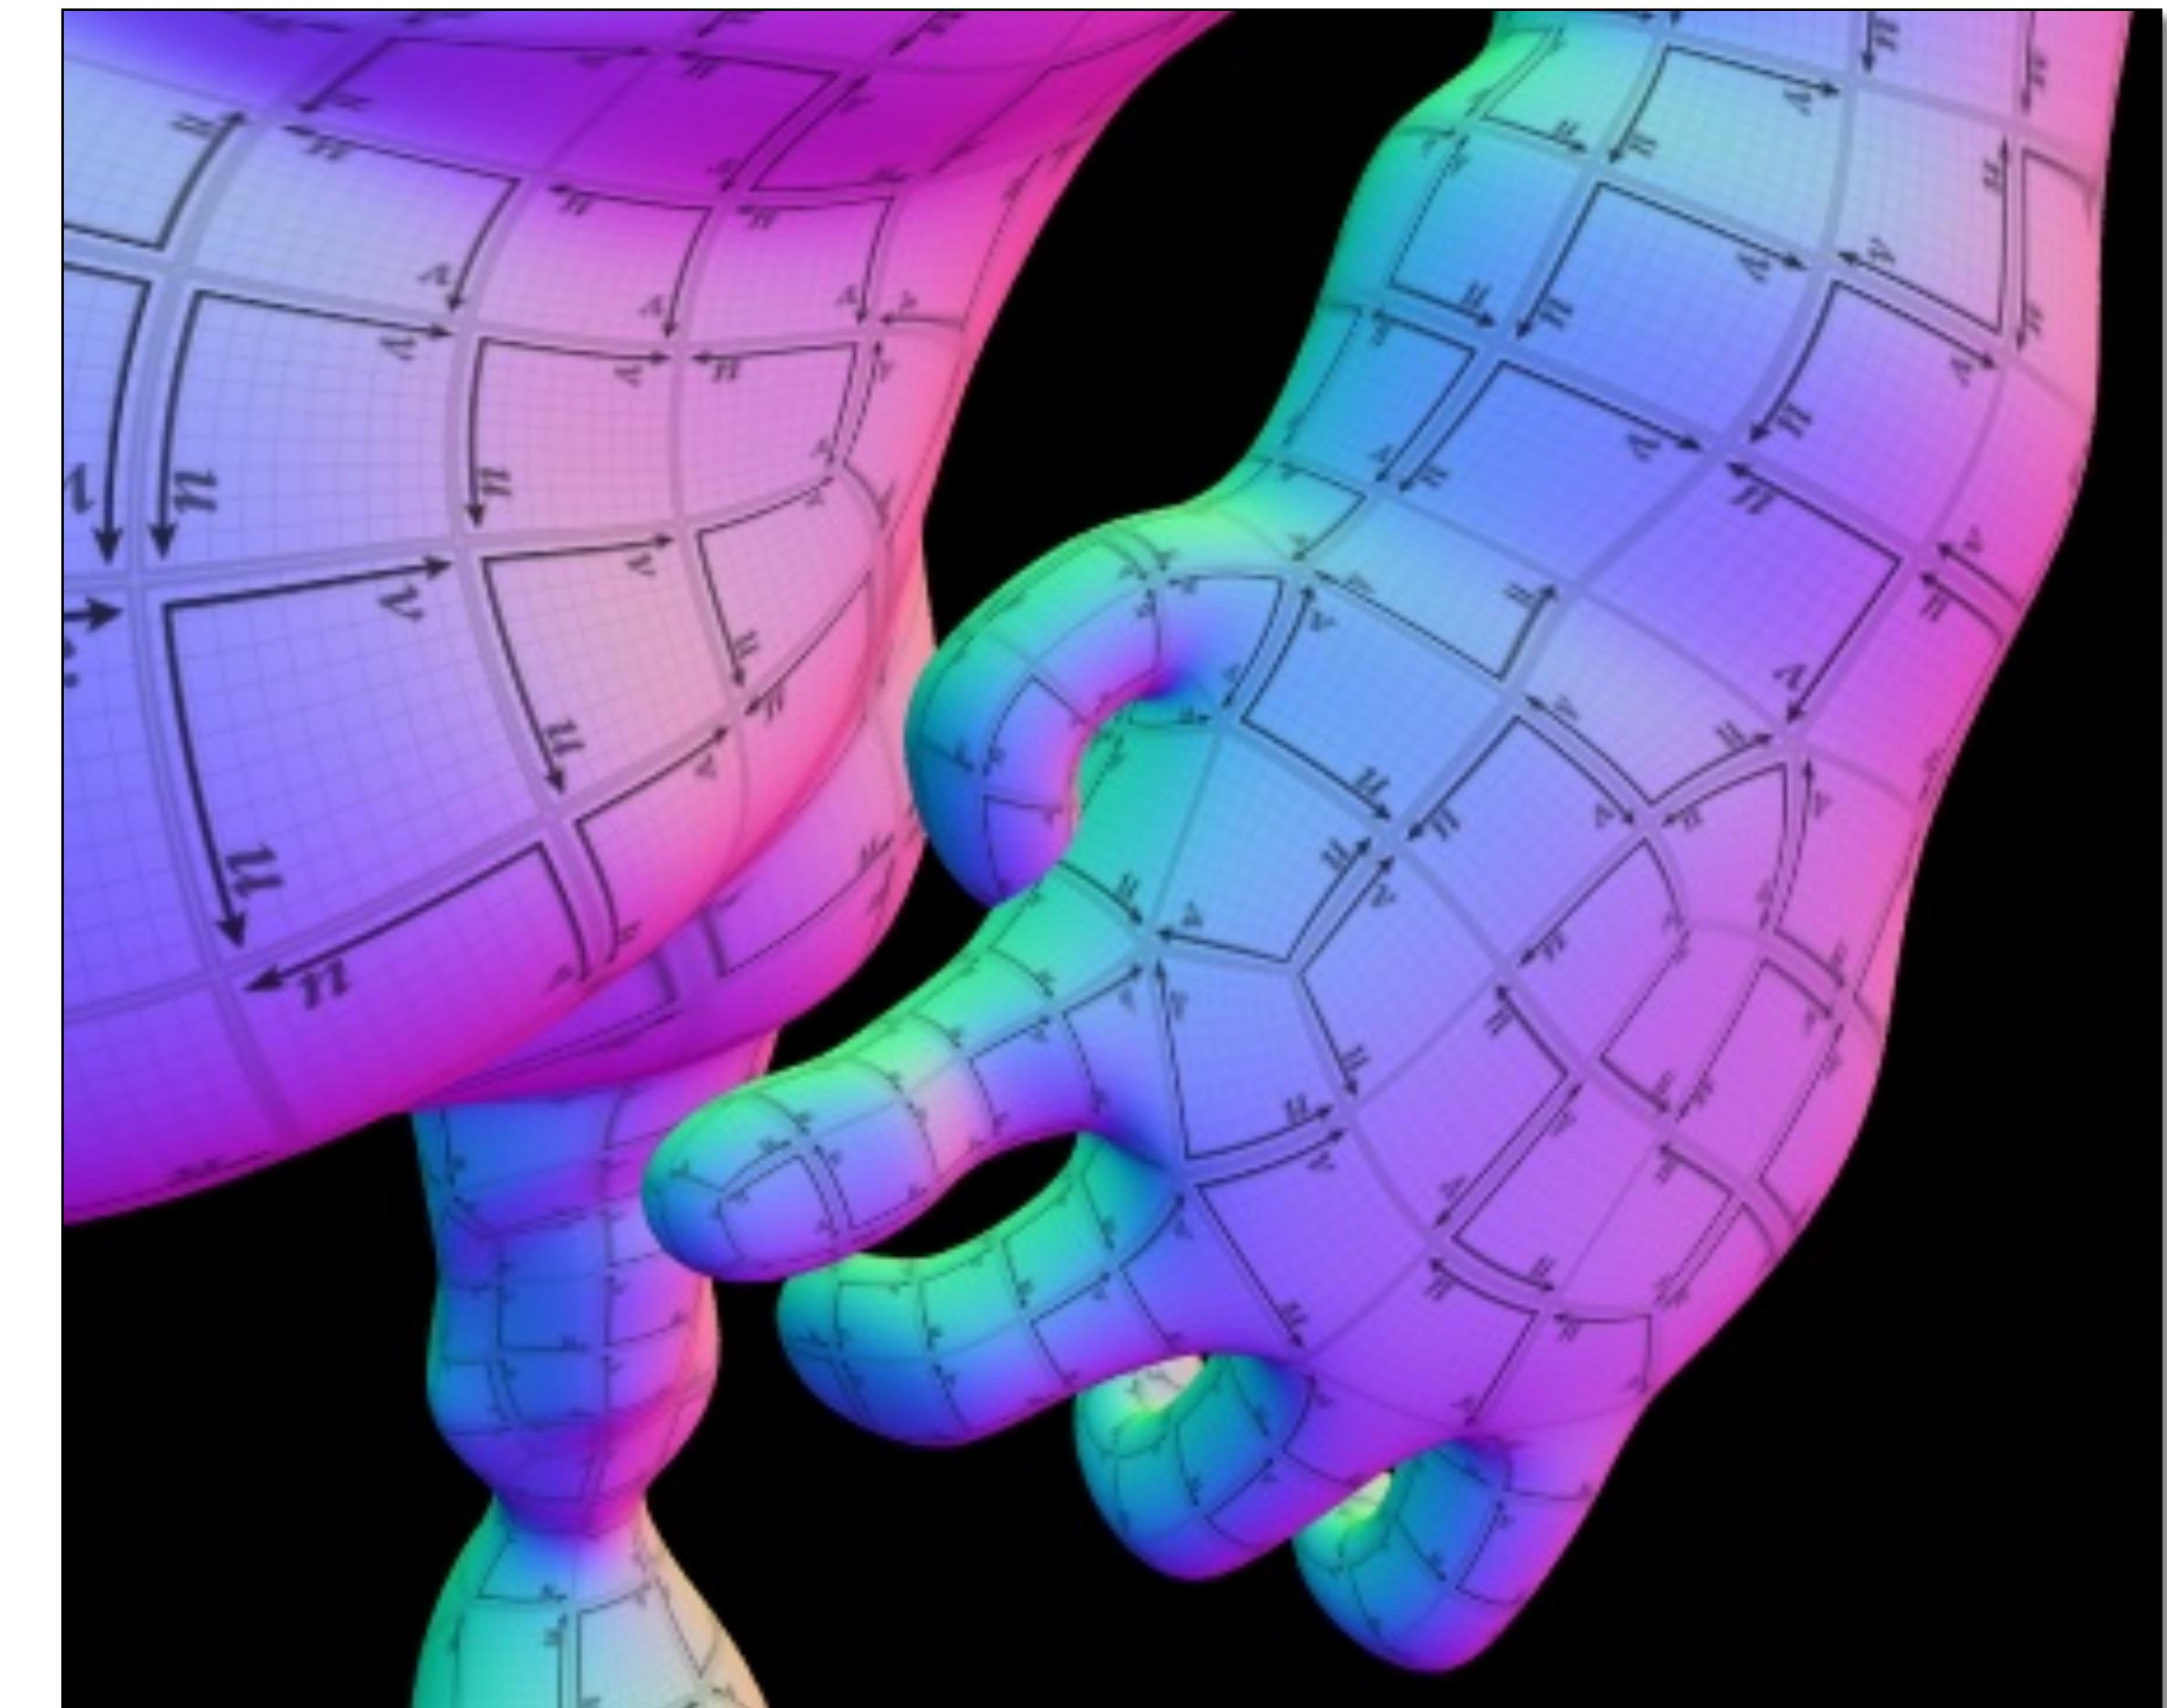
2D examples:
Moiré patterns, jaggies



Aliasing due to undersampling texture

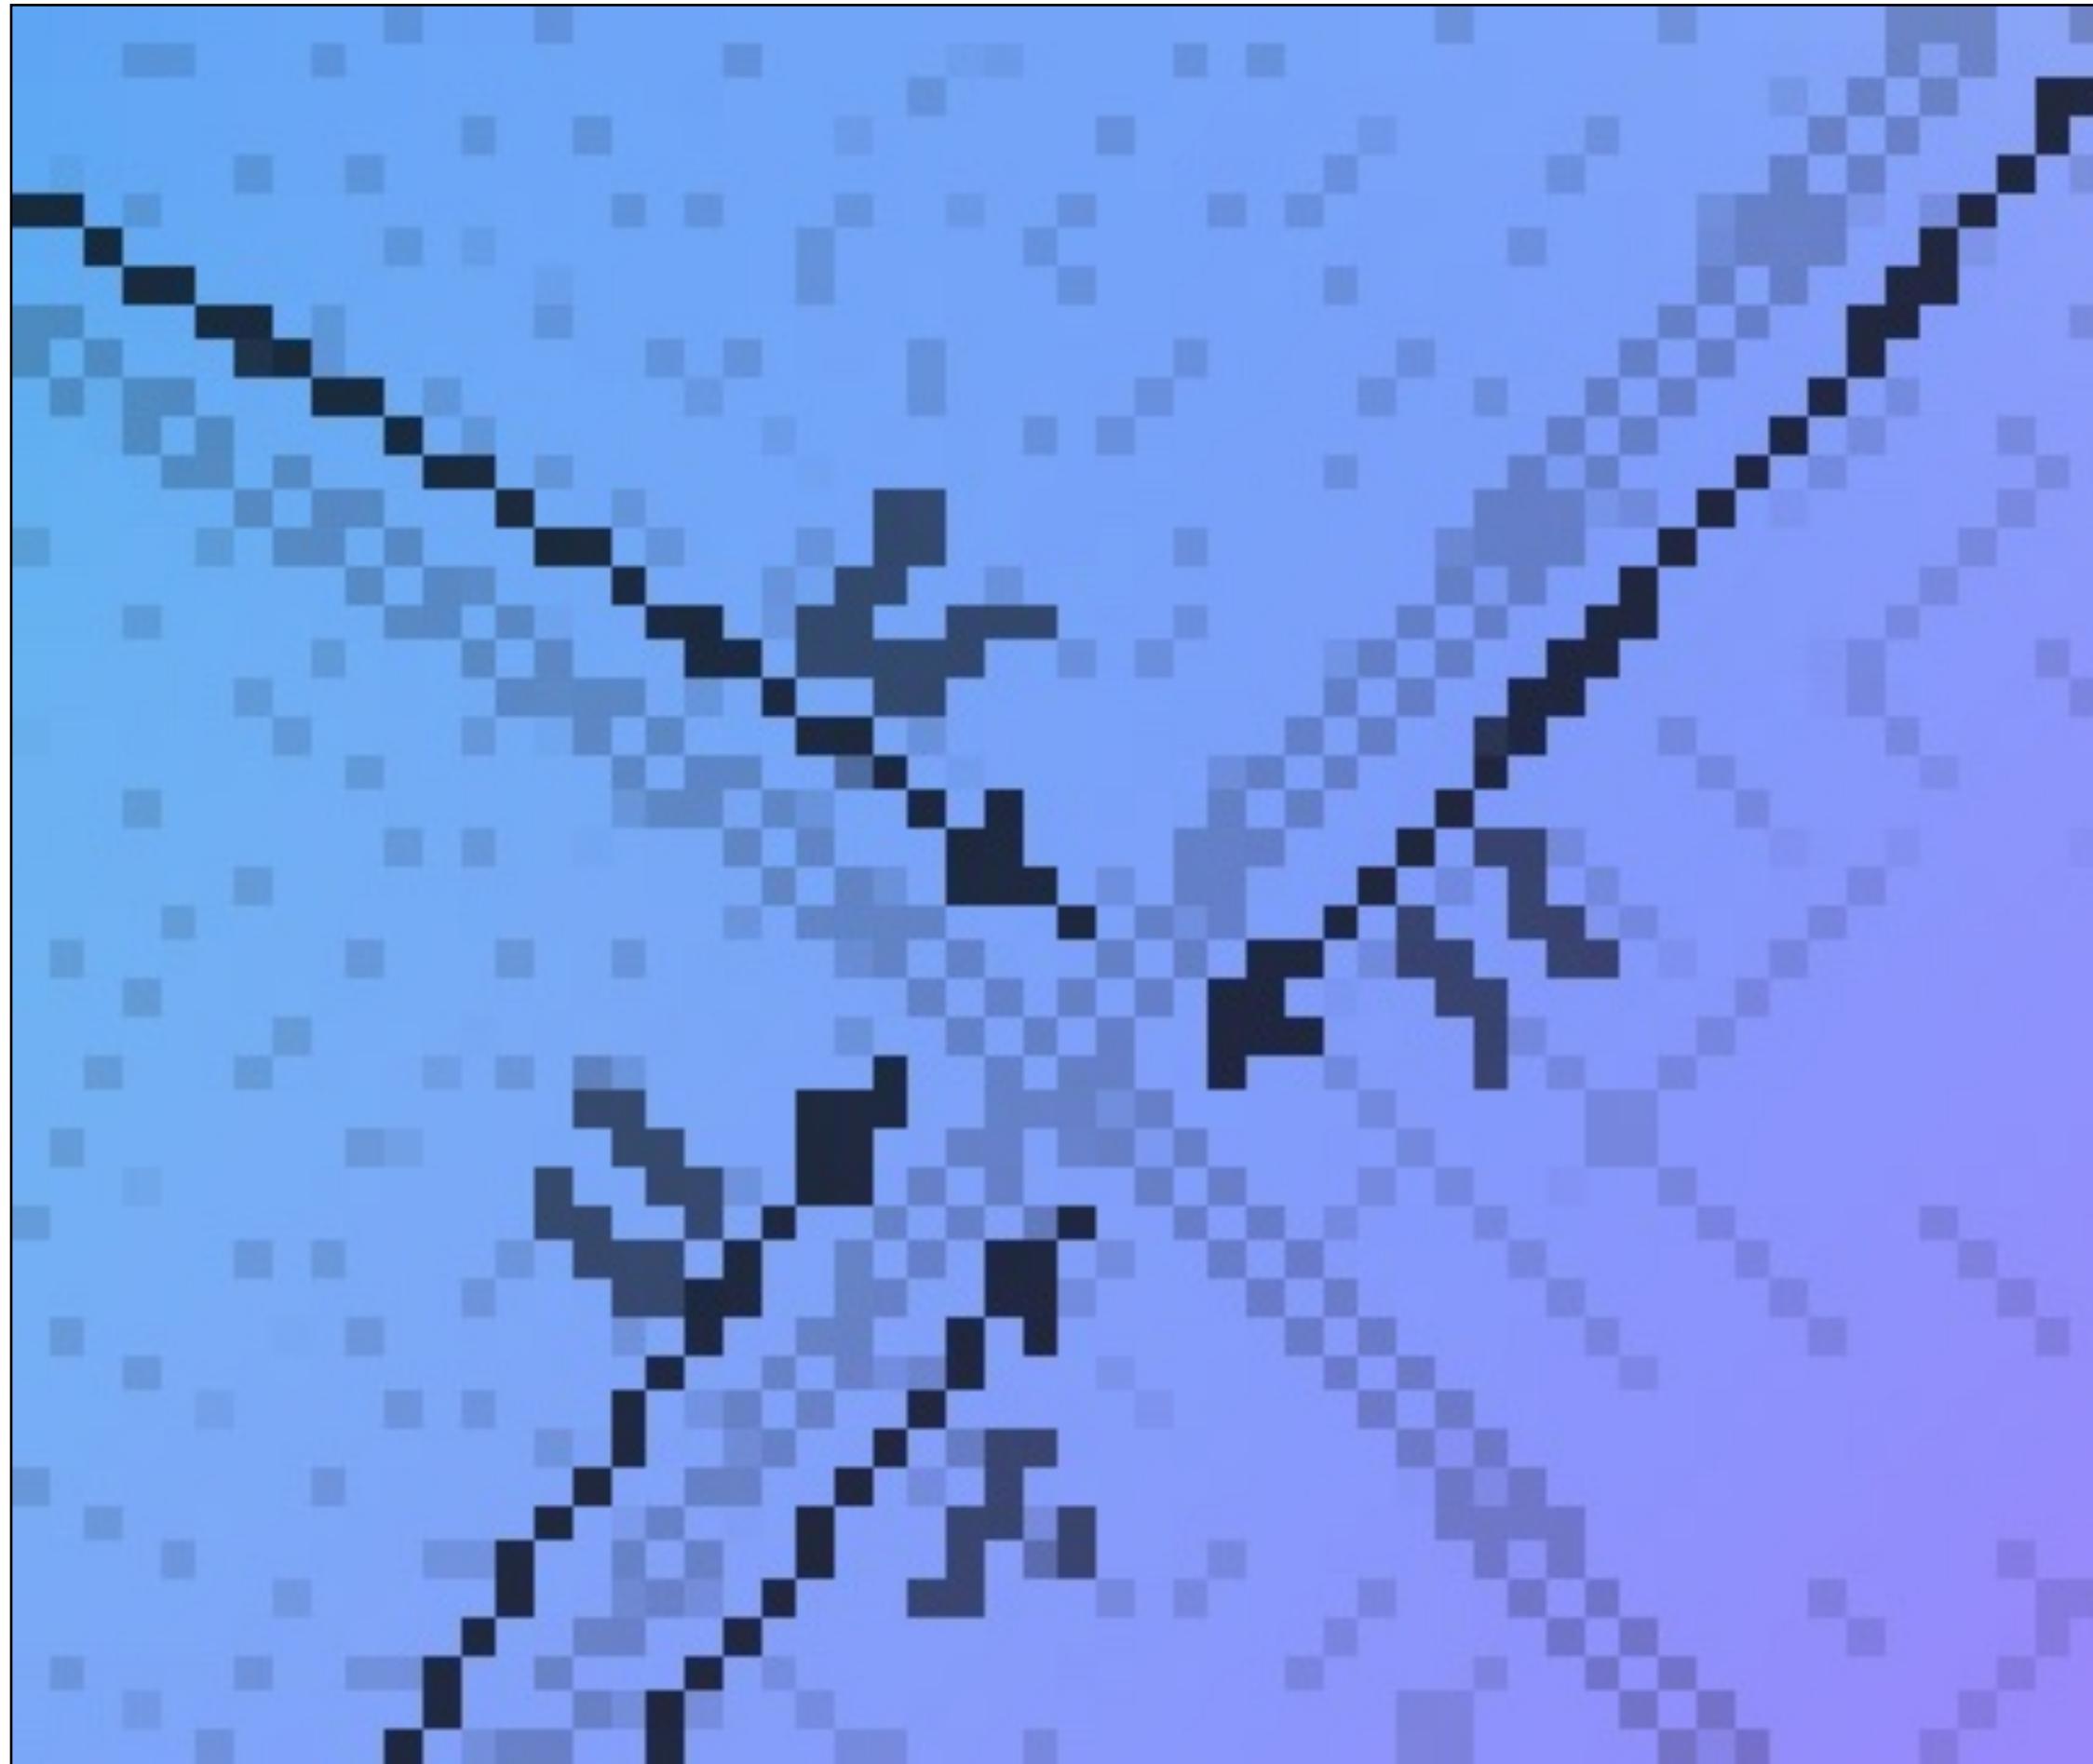


One texture sample per pixel
(aliasing!)



Anti-aliased texture sampling

Aliasing due to undersampling (zoom)

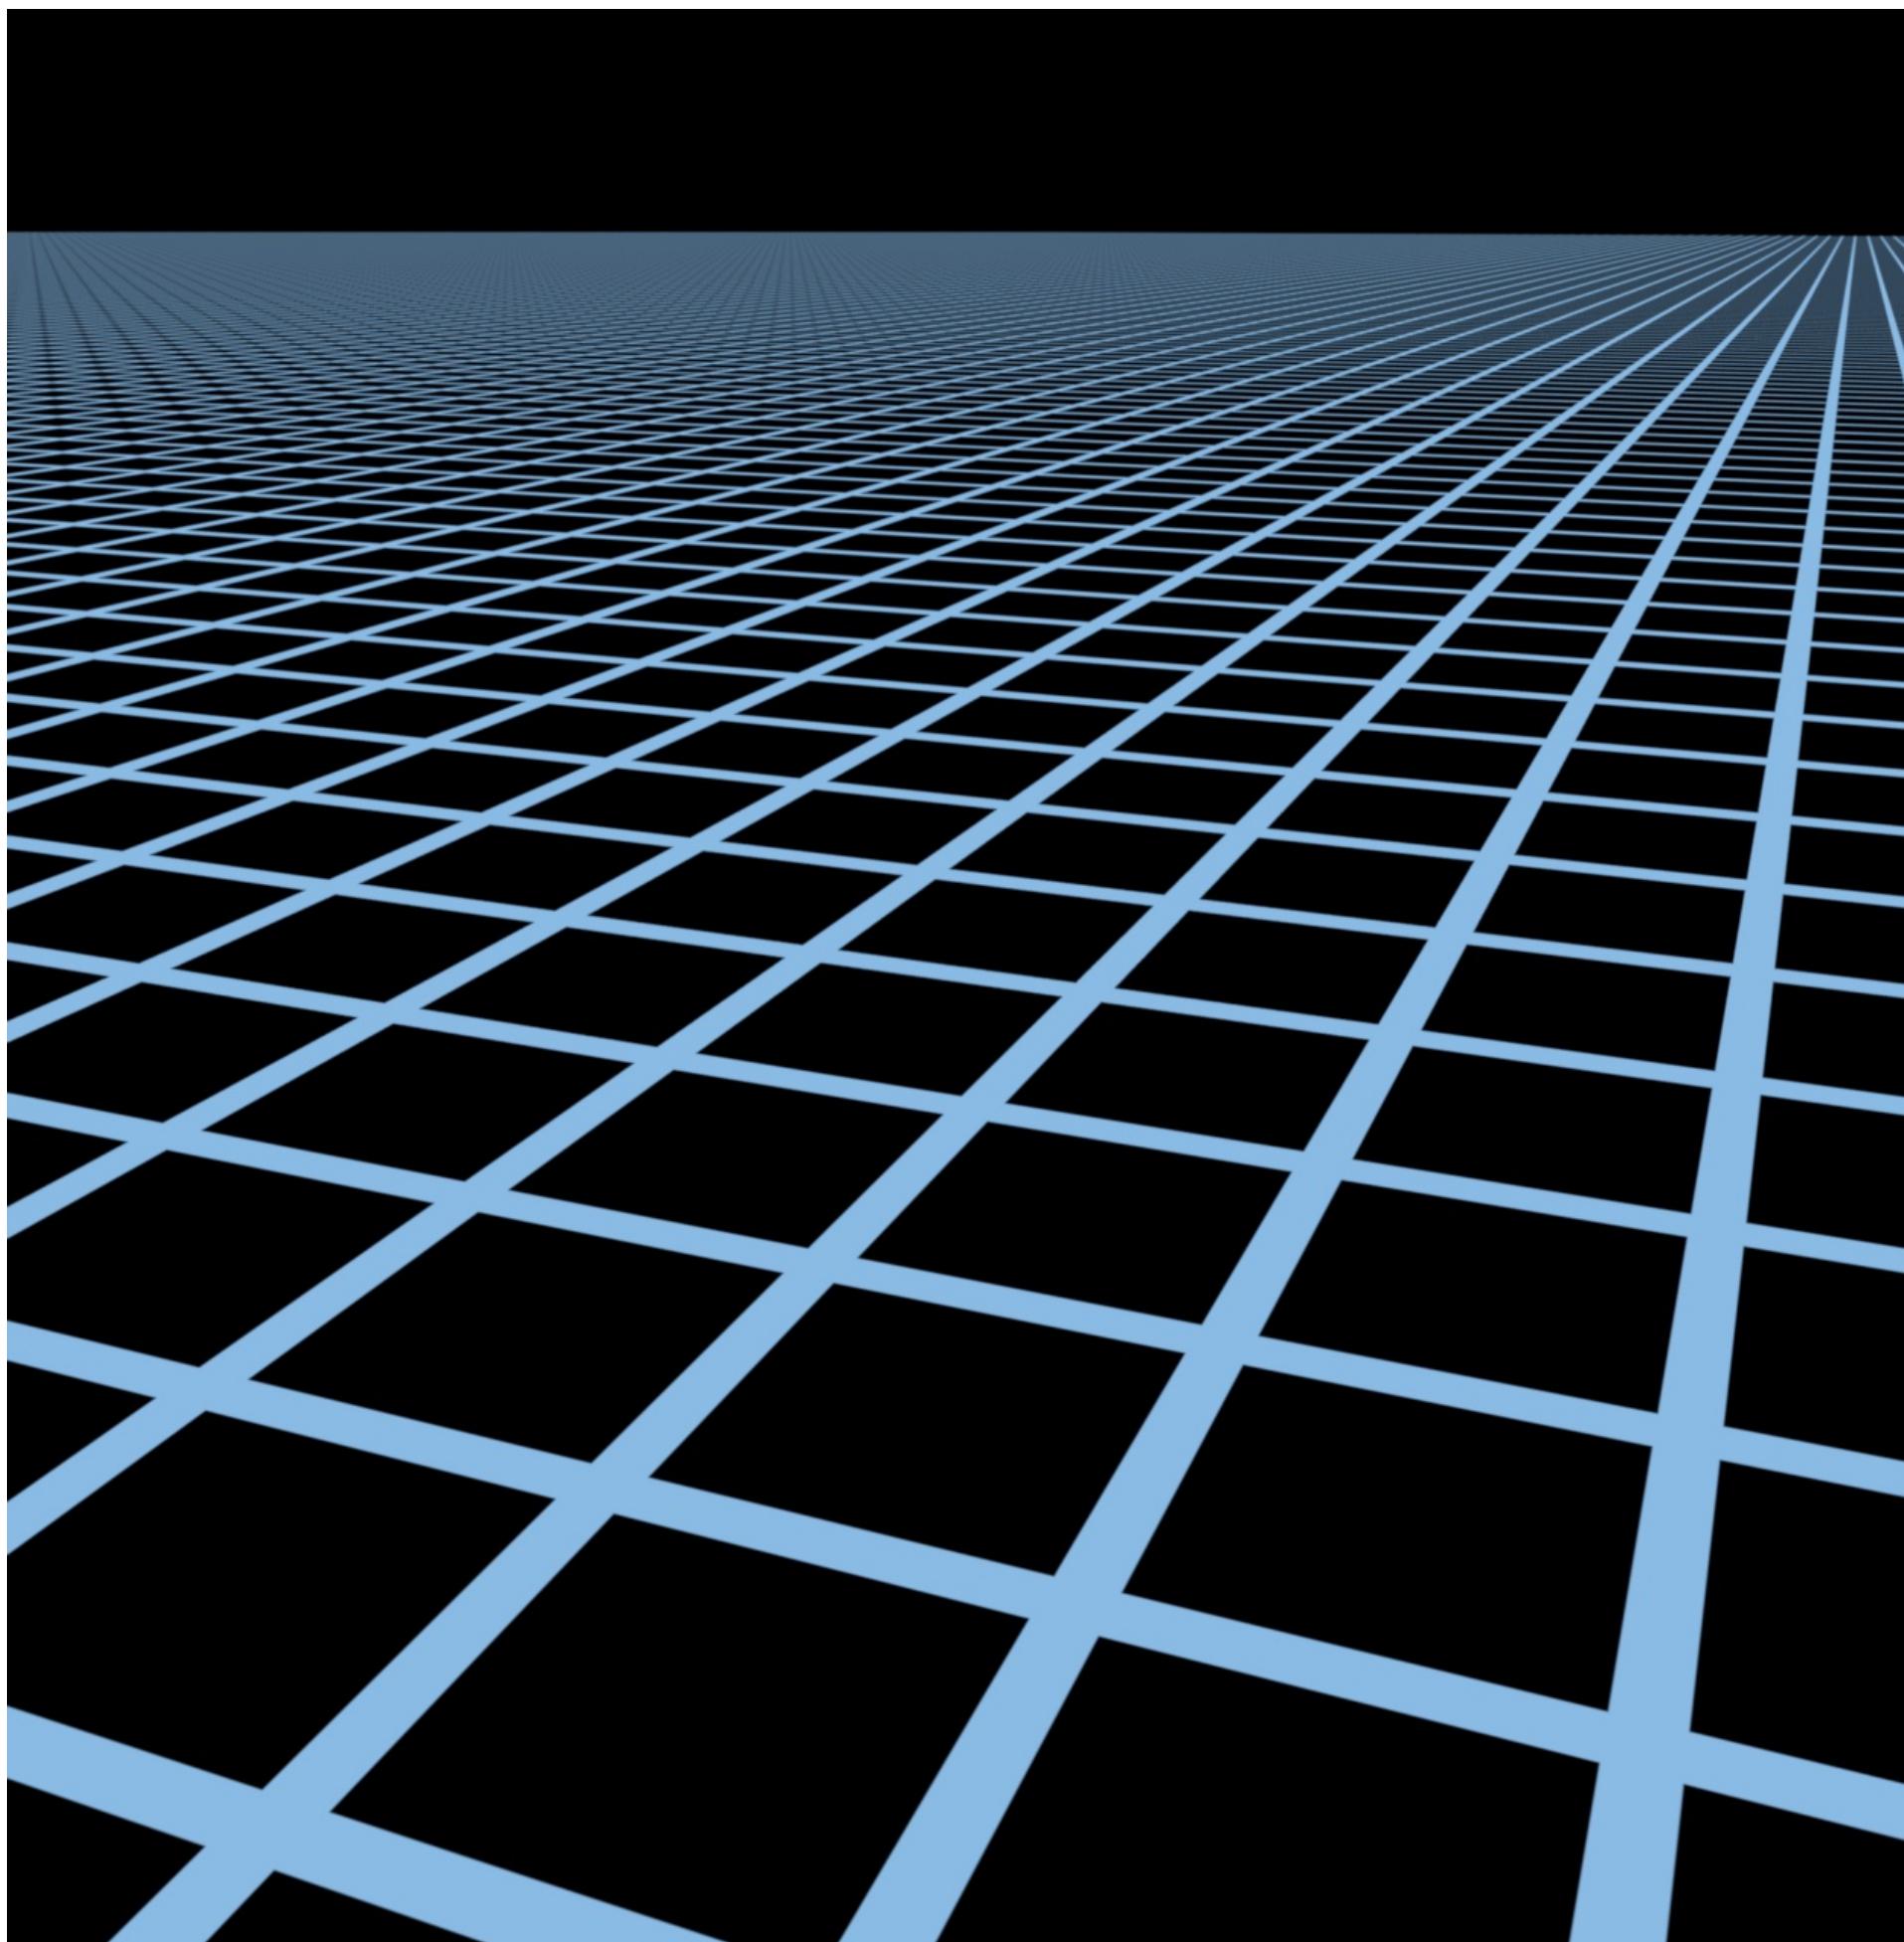


**One texture sample per pixel
(aliasing!)**

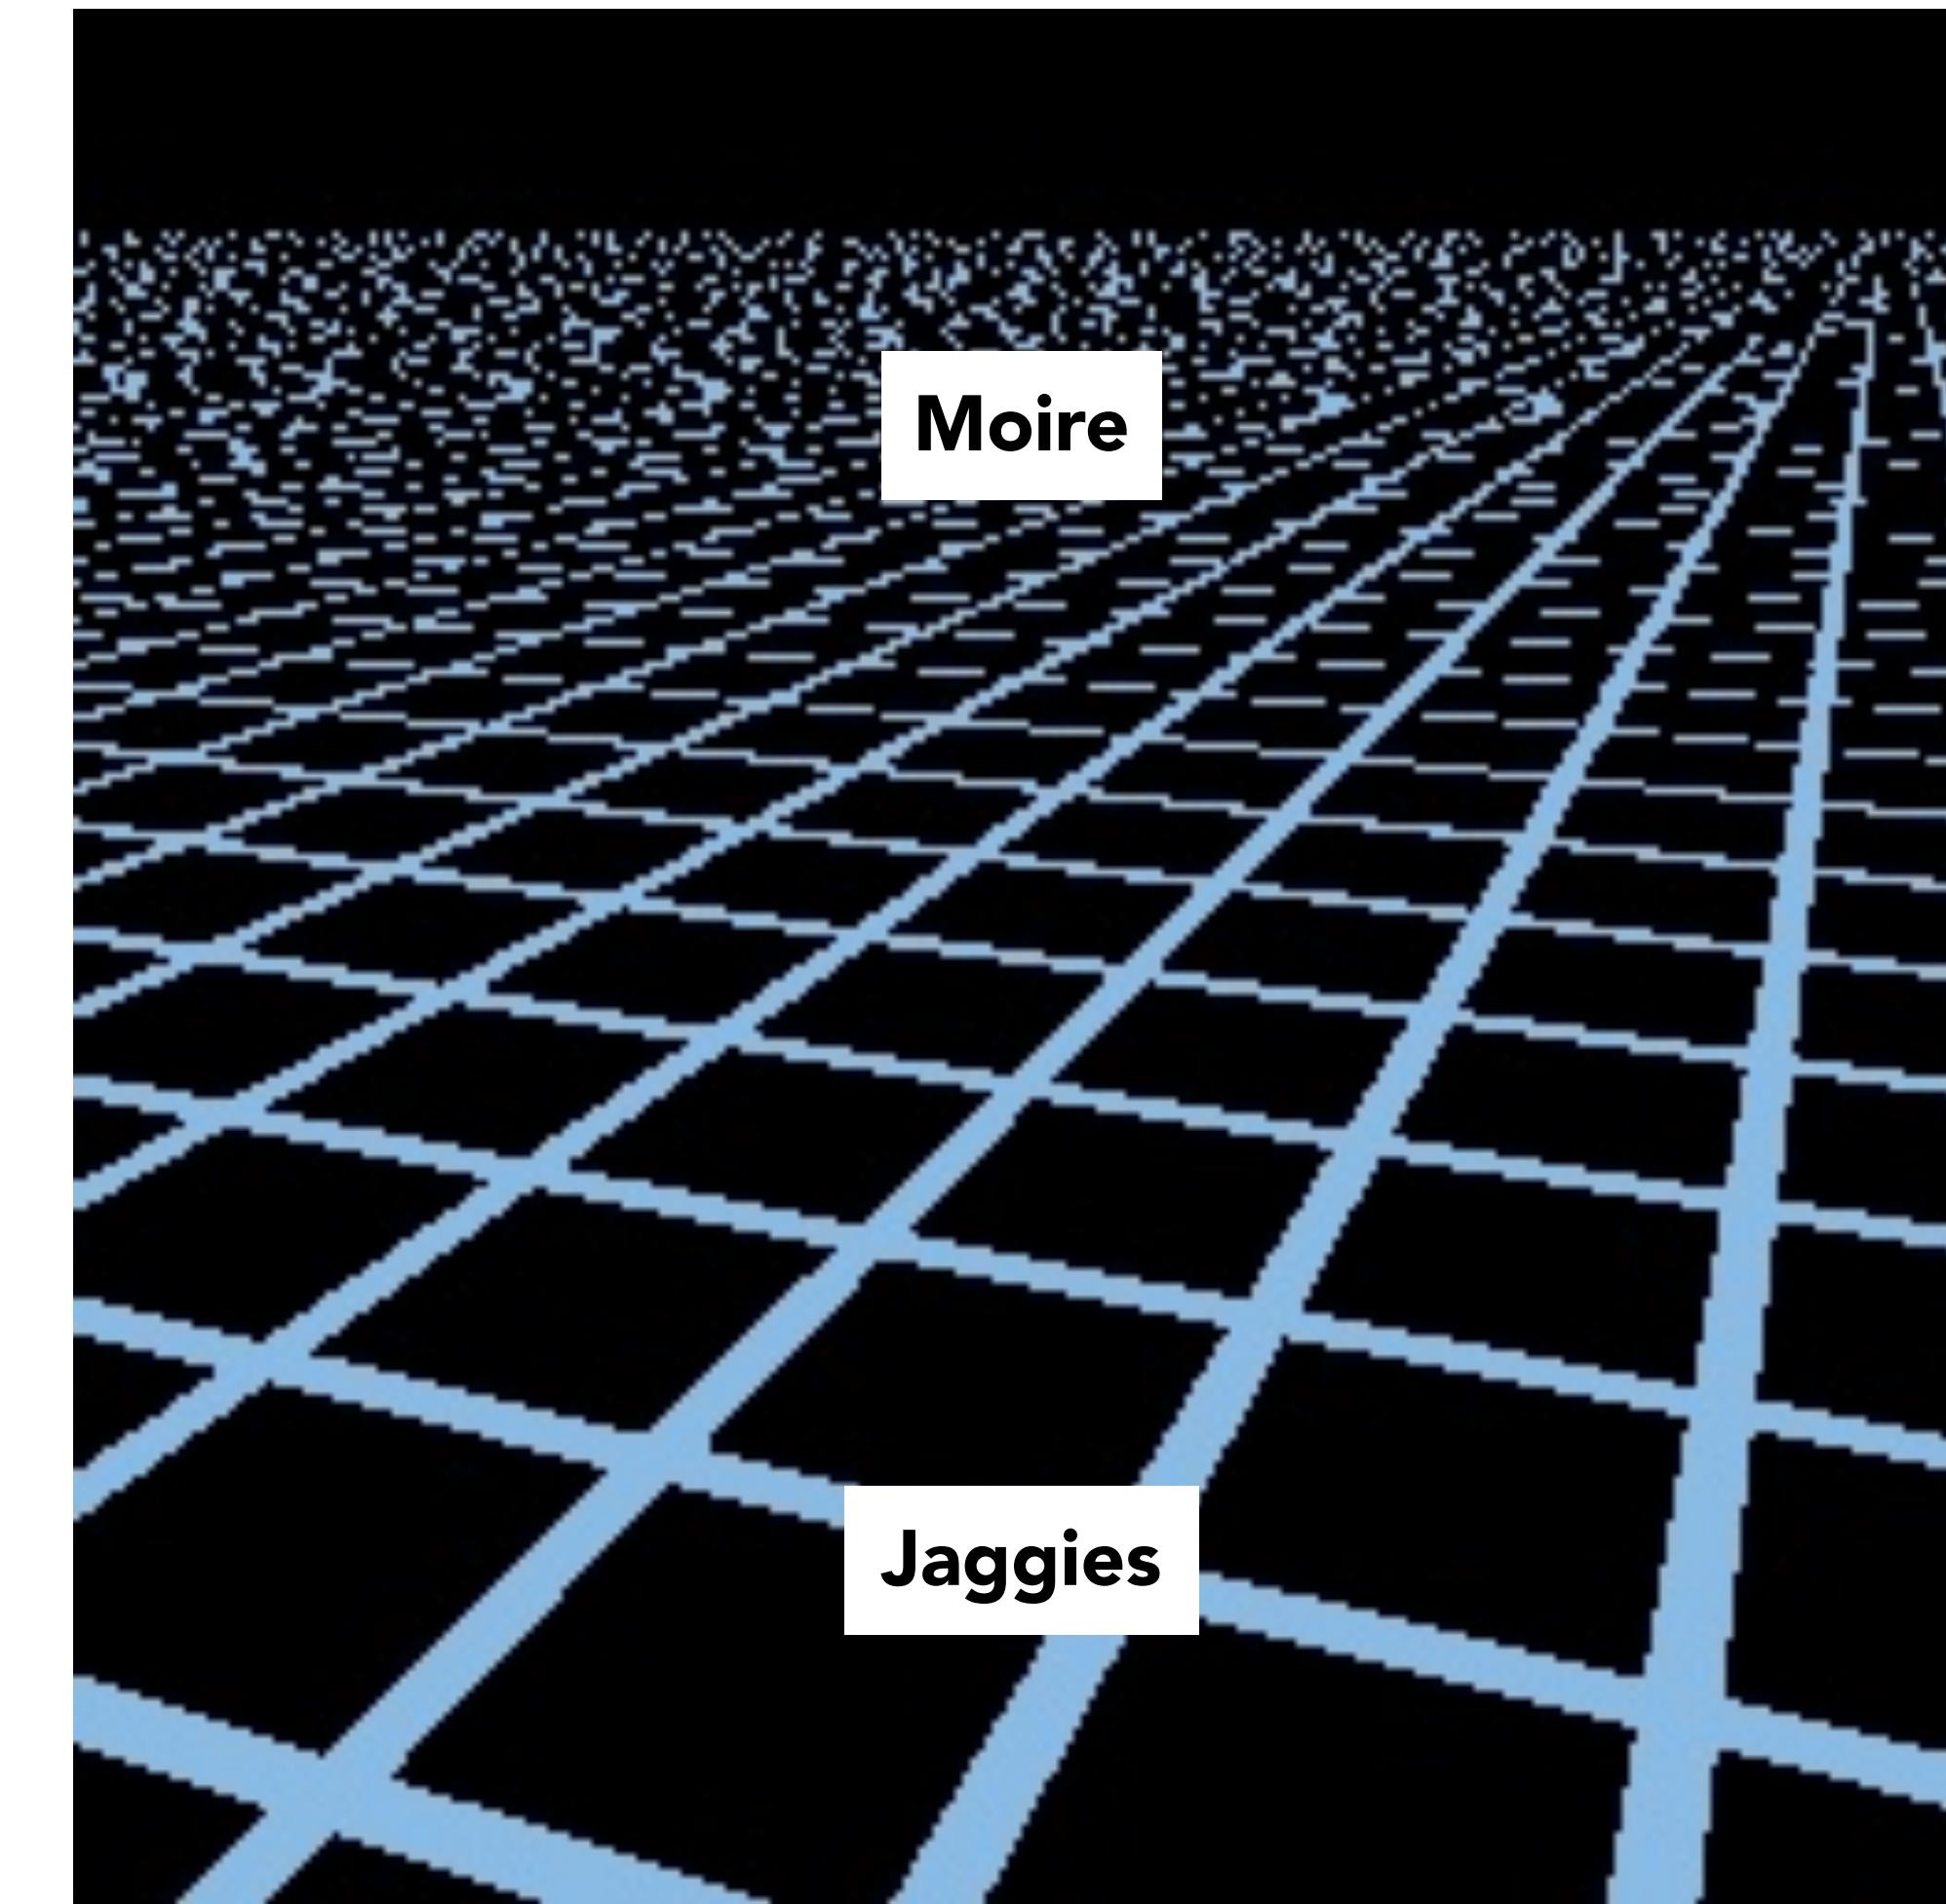


Anti-aliased texture sampling

Another example



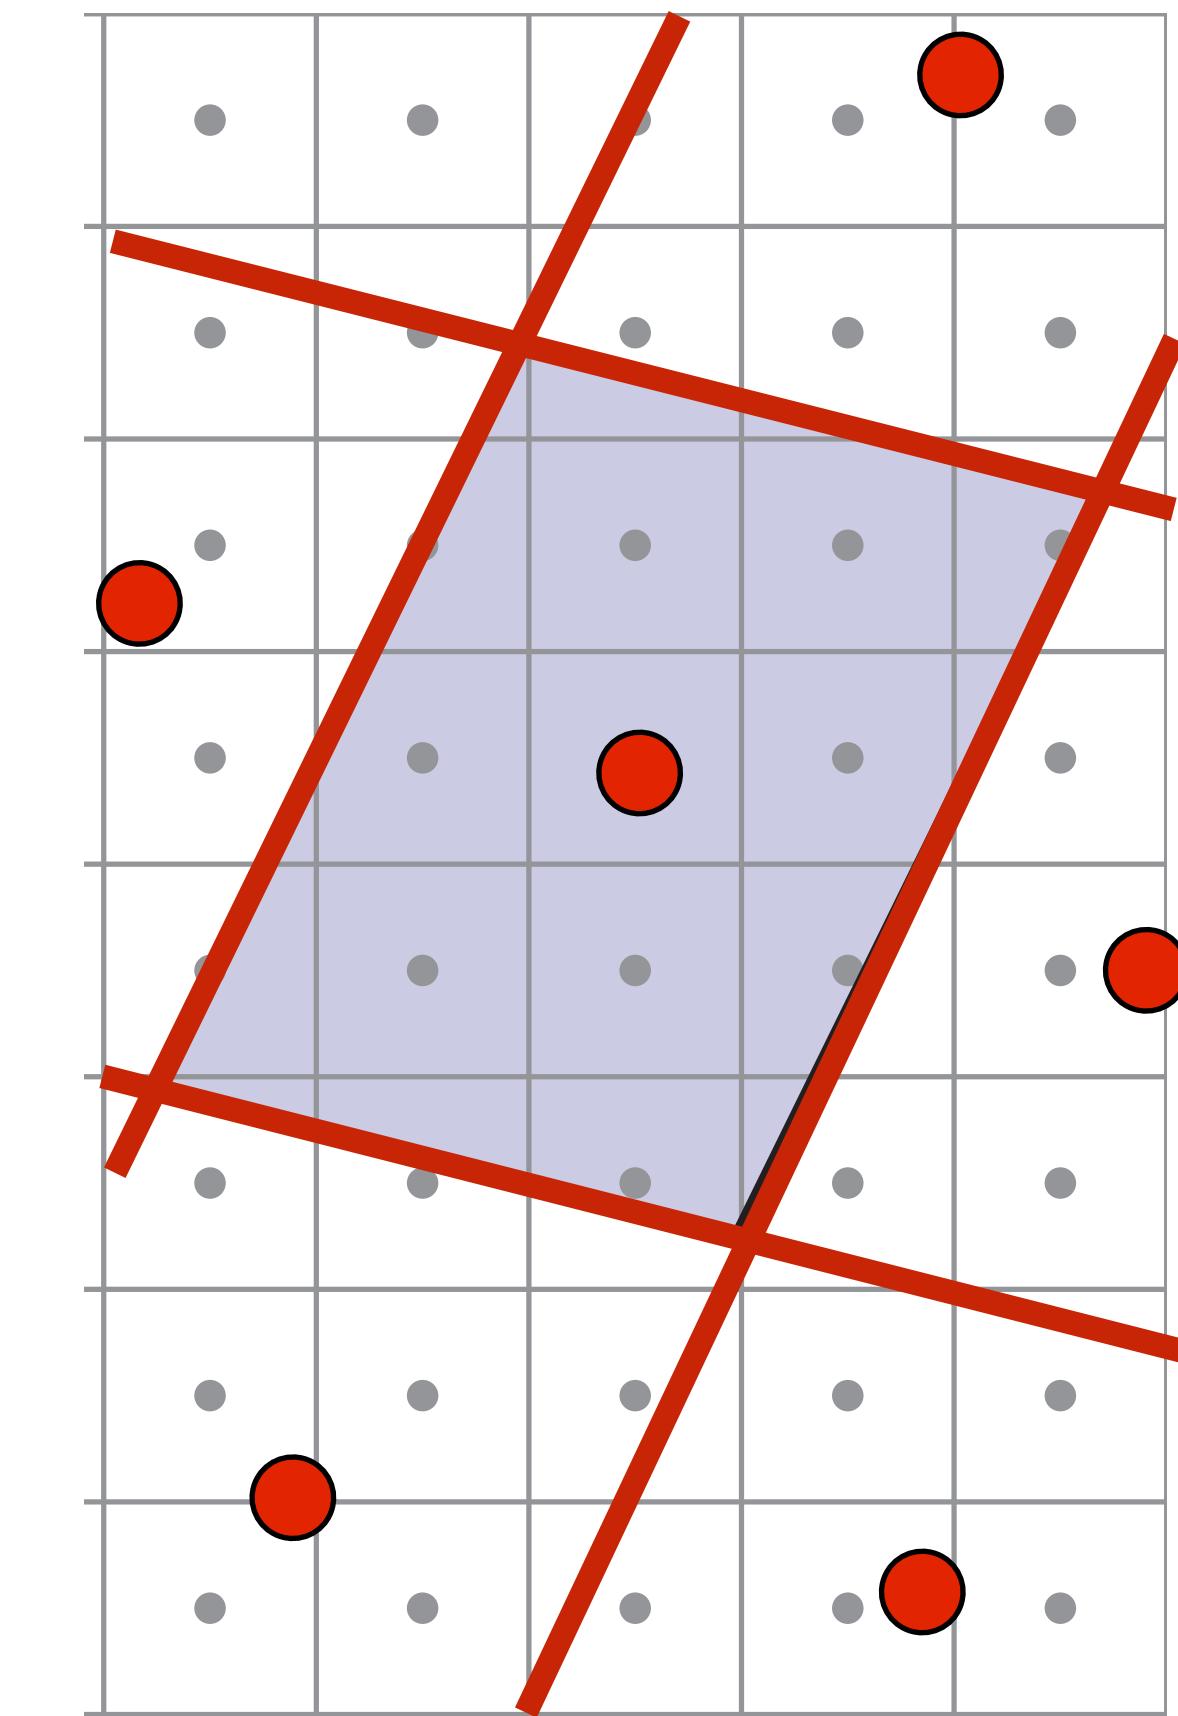
Anti-aliased result



Rendered image: 256x256 pixels

Texture minification - hard case

- **Challenge:**
 - **Many texels contribute to color of an output image pixel (sampling only one of them could yield aliasing)**
 - **Shape of pixel footprint can be complex**

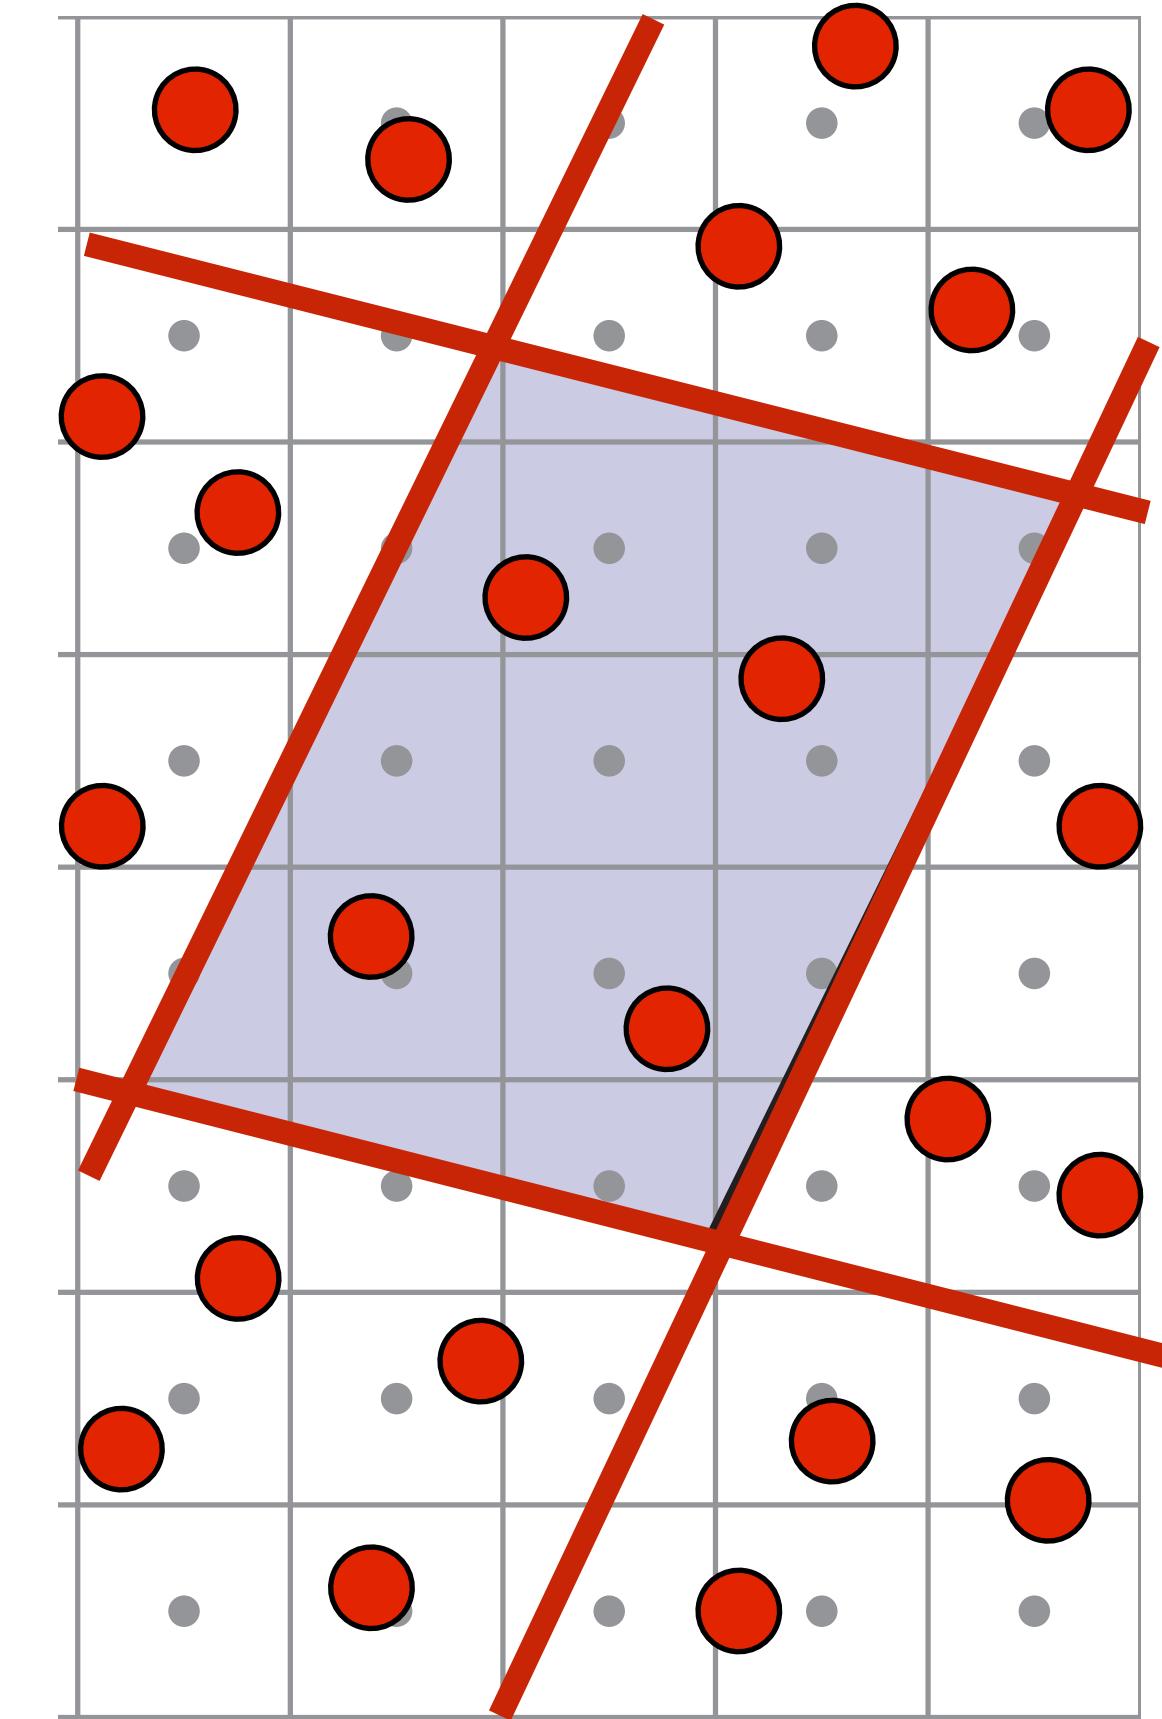


Shaded region = pixel area
Red lines = screen pixel boundaries
Red dots = texture space sample points for adjacent pixels

Texture minification - hard case

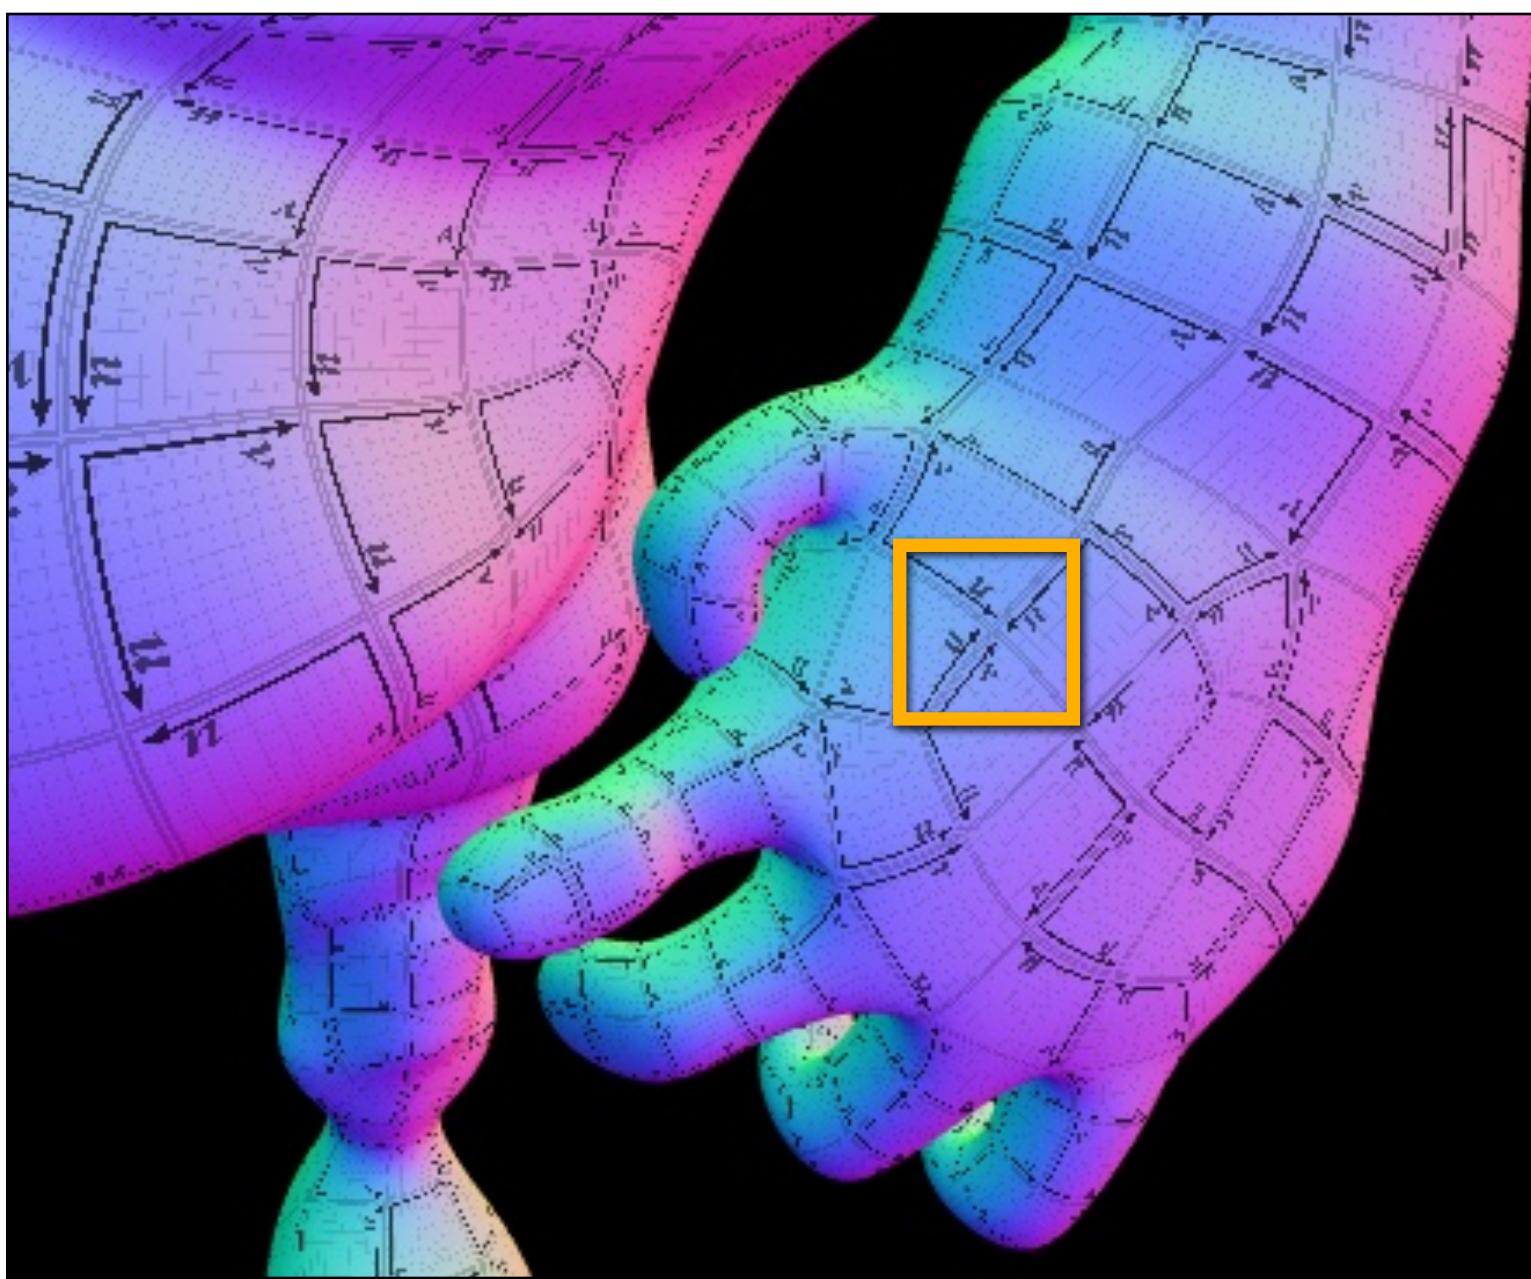
- **Challenge:**
 - Many texels contribute to color of an output image pixel (sampling only one of them could yield aliasing)
 - Shape of pixel footprint can be complex
- **One solution that you already know: supersampling**
 - Averaging many texture samples per pixel can approximate result of convolving texture map with pixel-area sized filter
 - Problem?

Alternative solution: remove high frequency from texture to reduce aliasing!

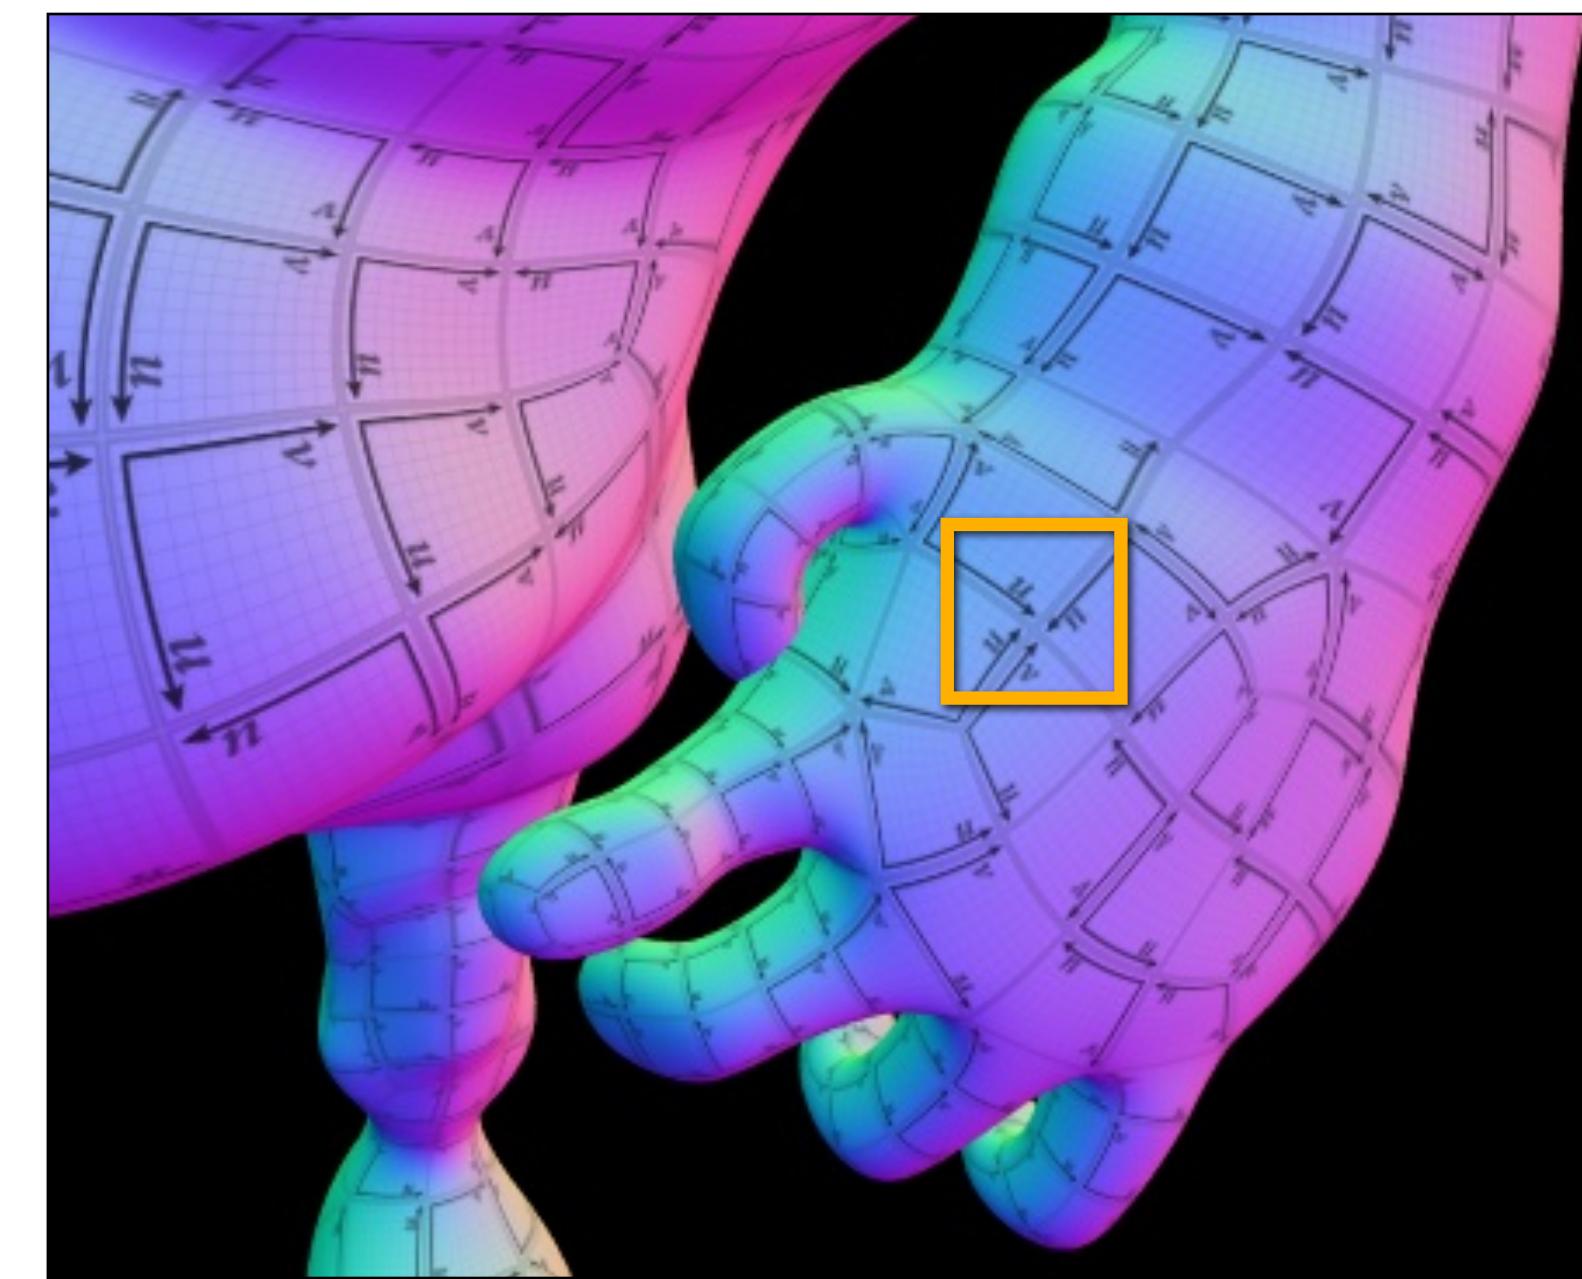
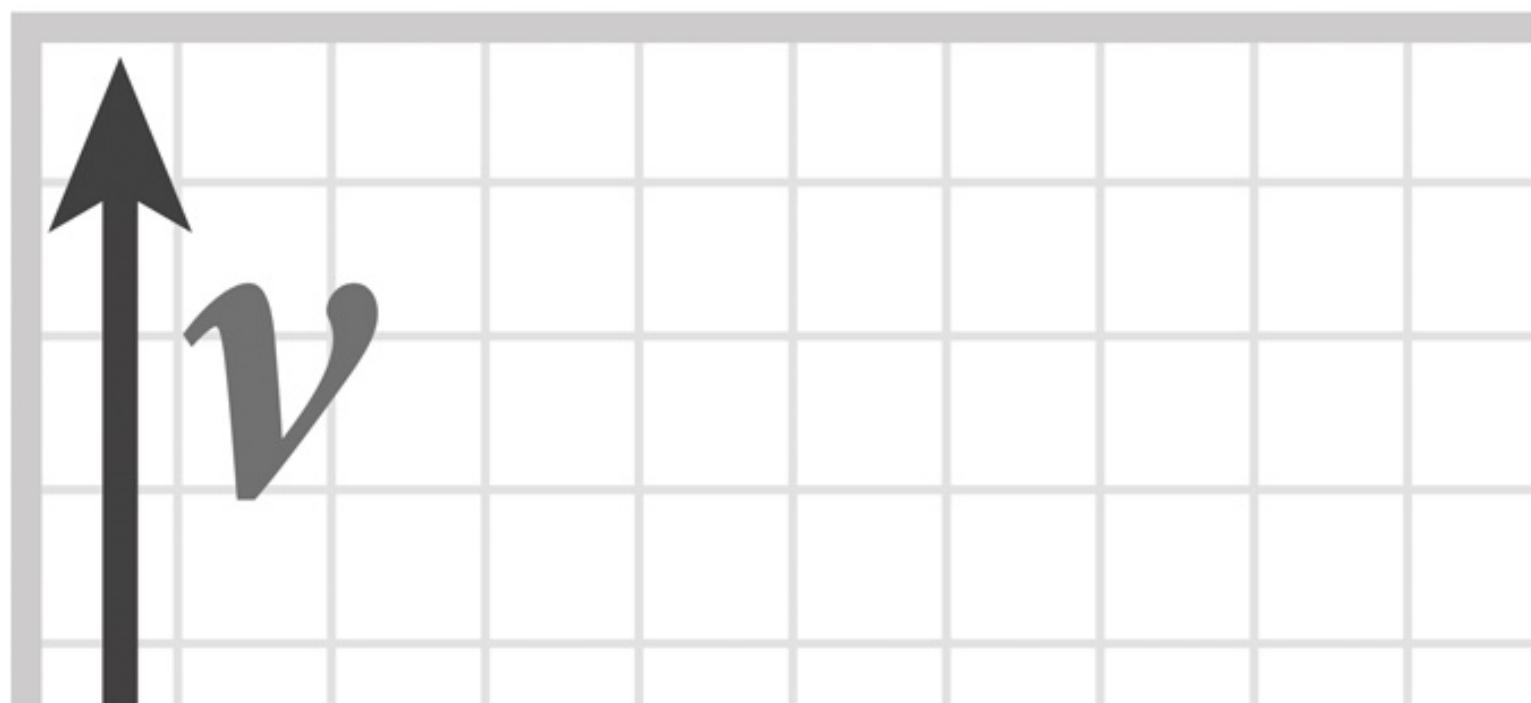


Shaded region = pixel area
Red lines = screen pixel boundaries
Red dots = texture space sample points for adjacent pixels

Pre-filtering texture map reduces aliasing



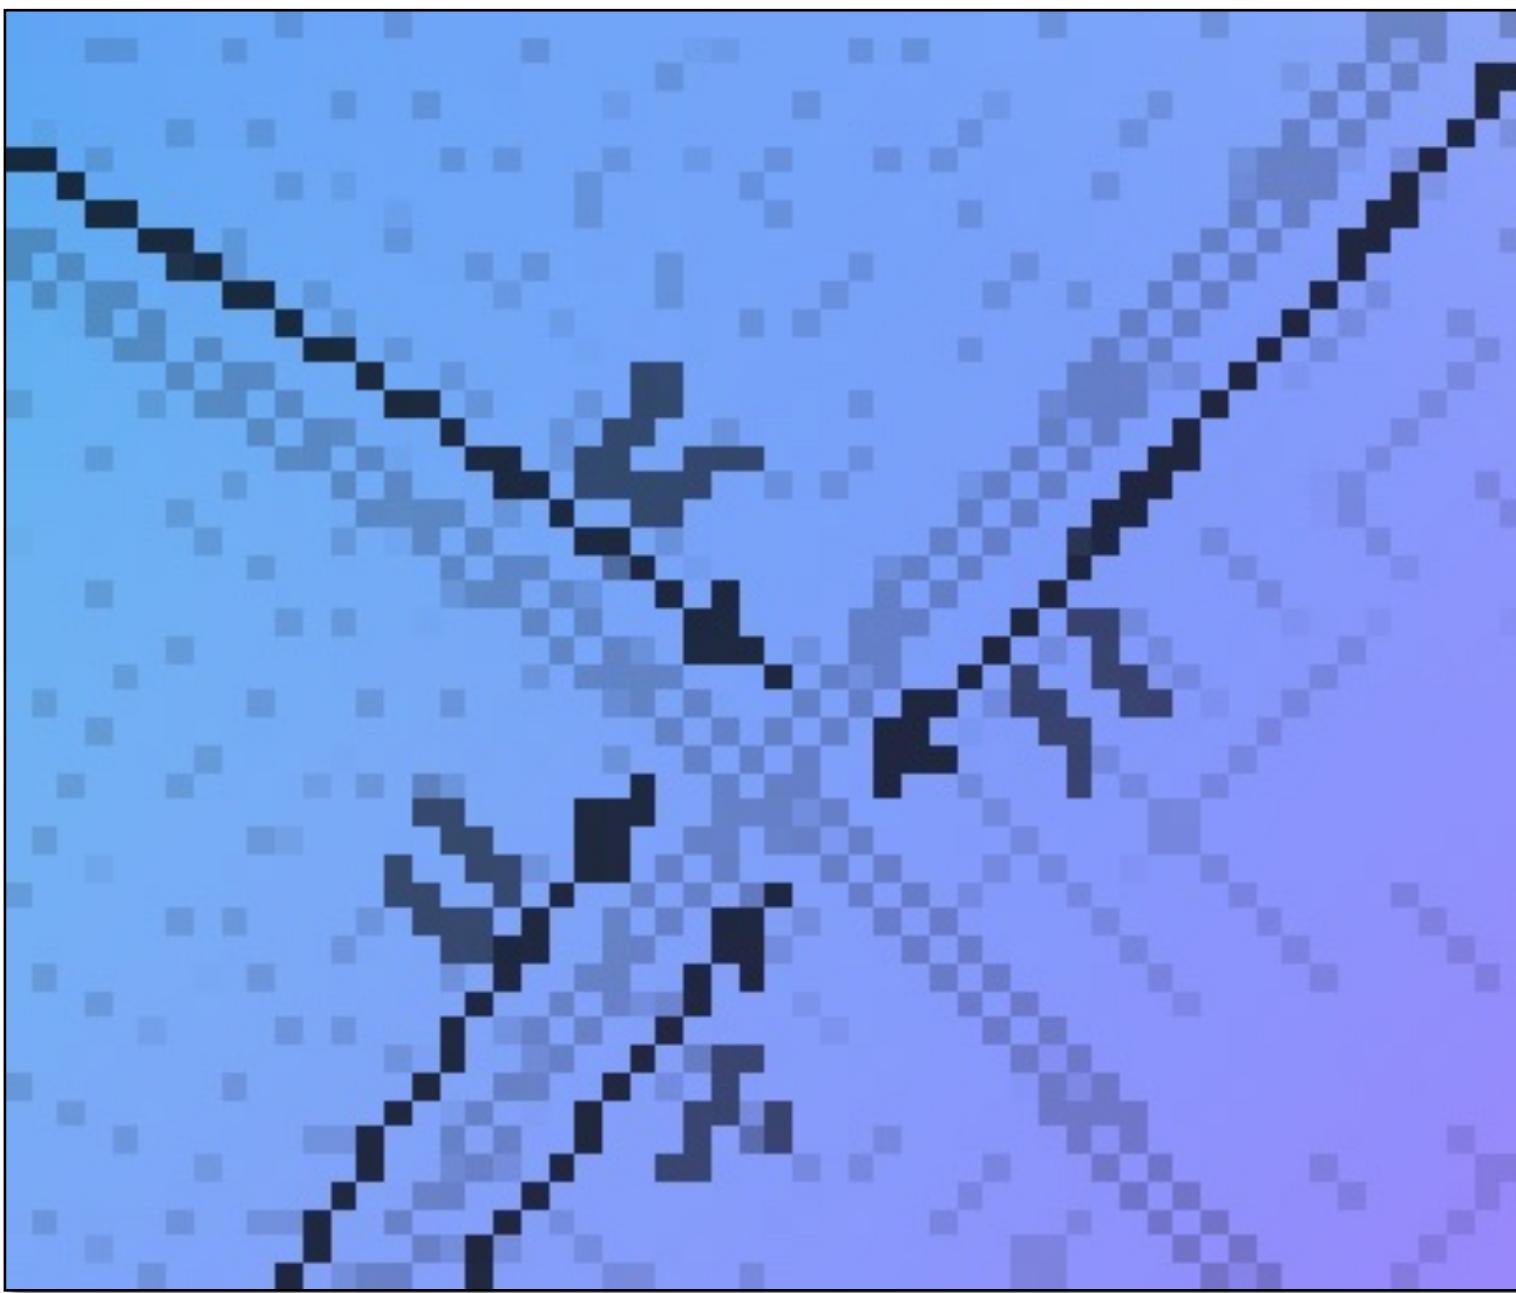
One texture sample per pixel
(aliasing!)



Pre-filtered texture map
(high frequencies removed)



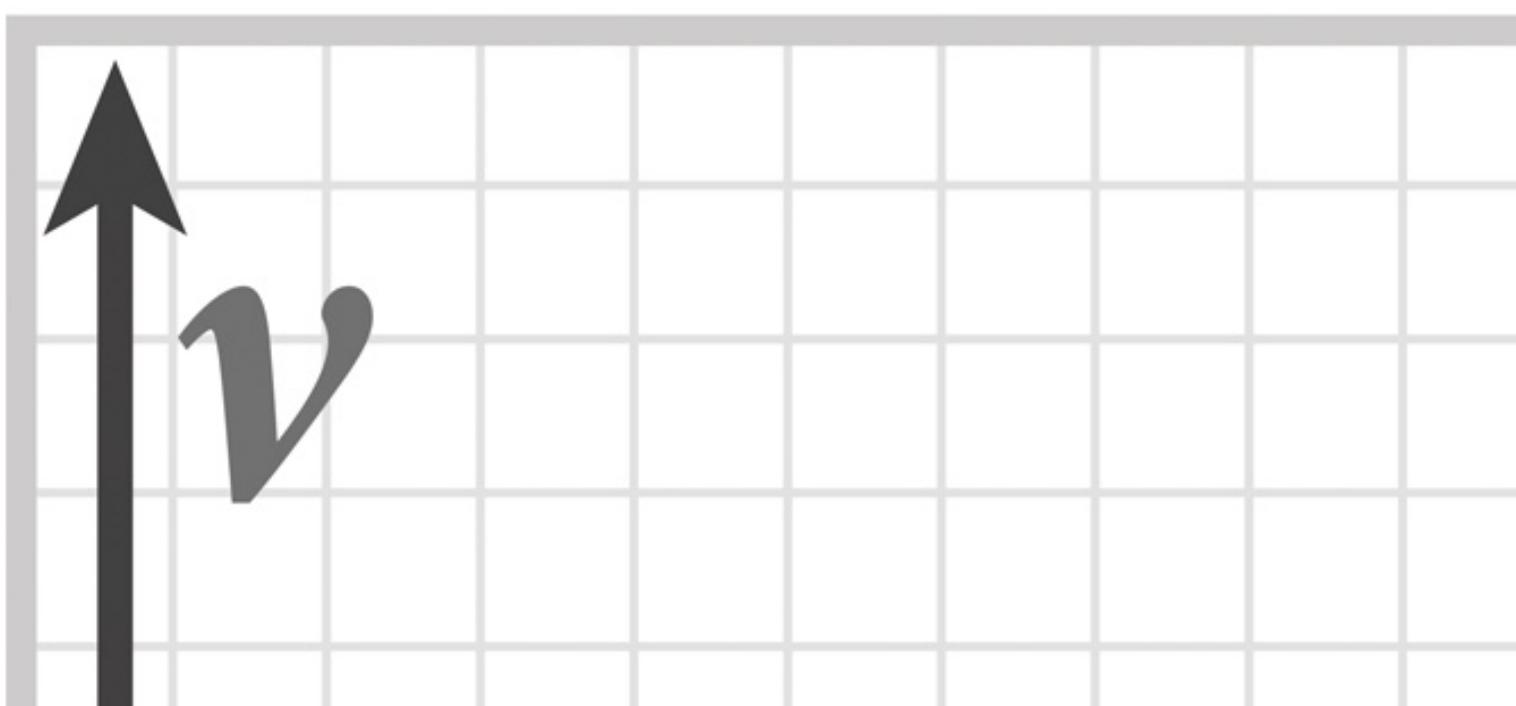
Pre-filtering texture map reduces aliasing



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

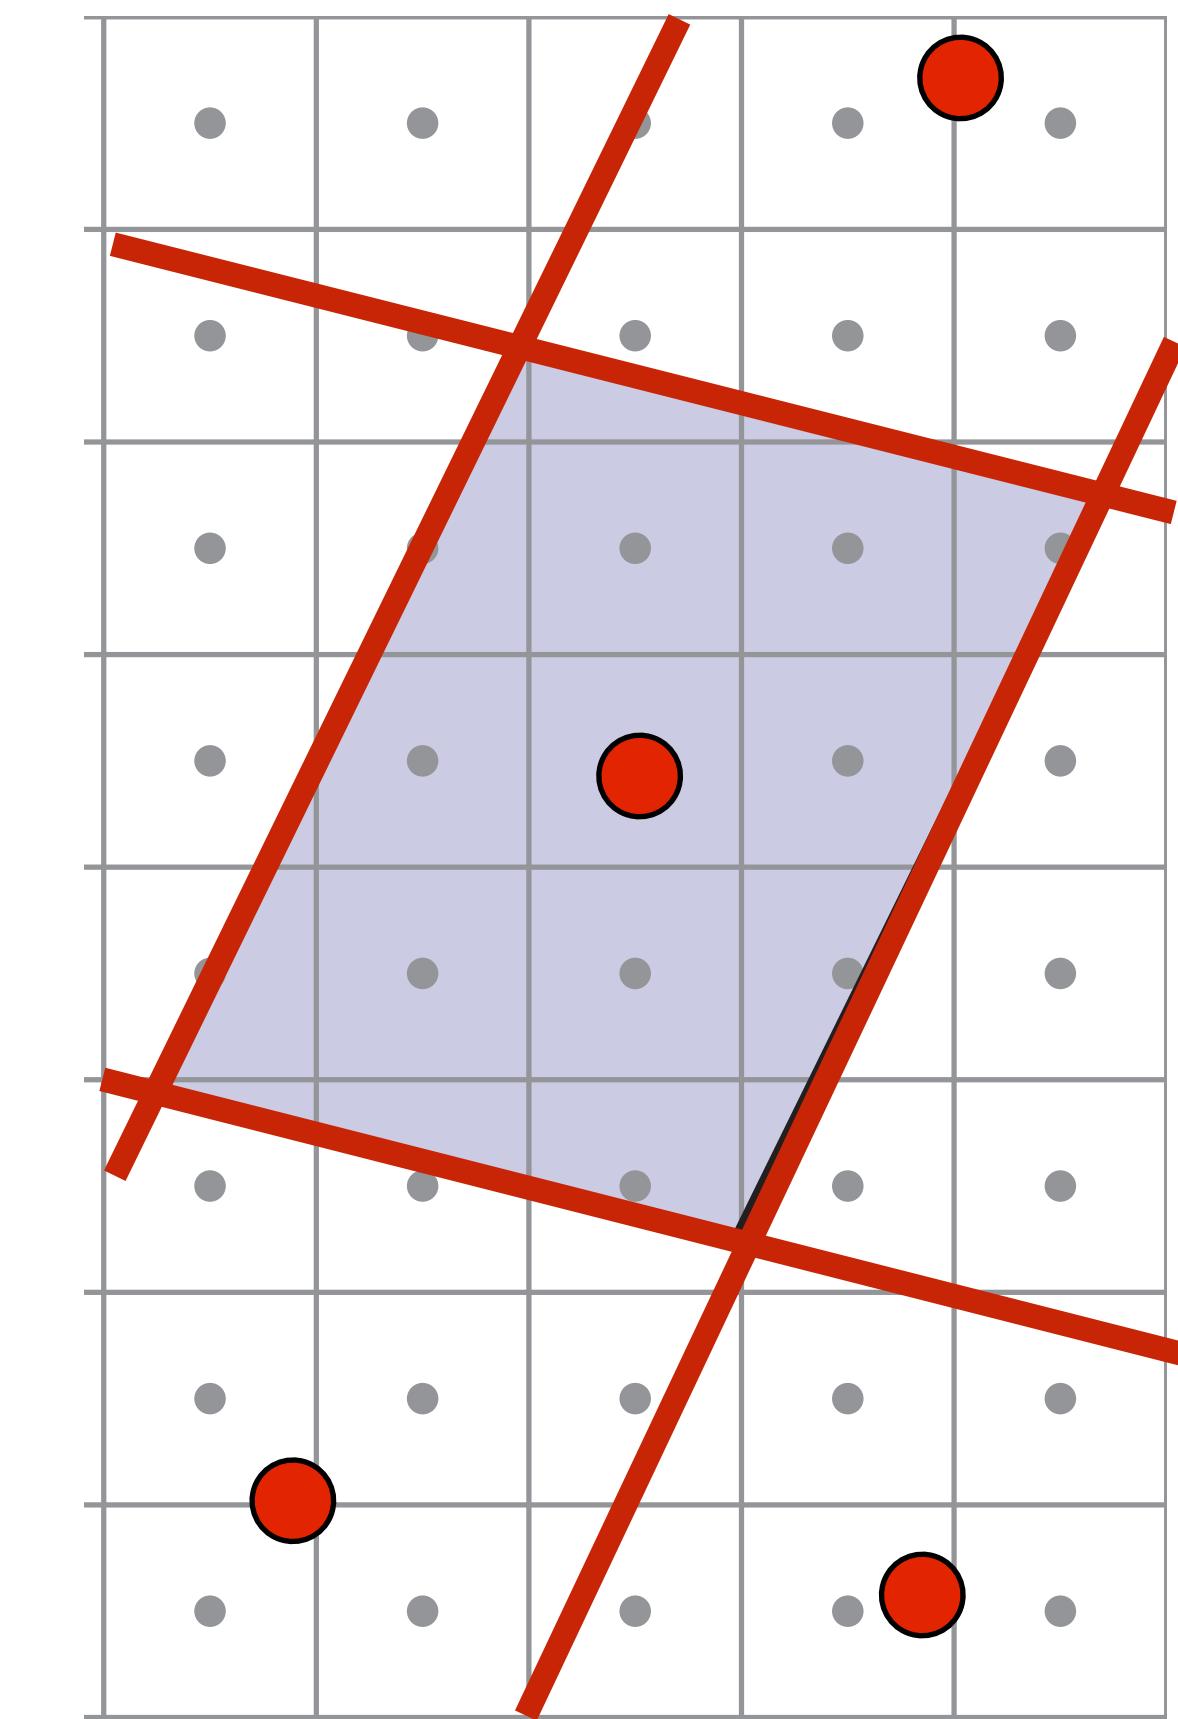


Pre-filtered texture map
(high frequencies removed)



But how much should we pre-filter?

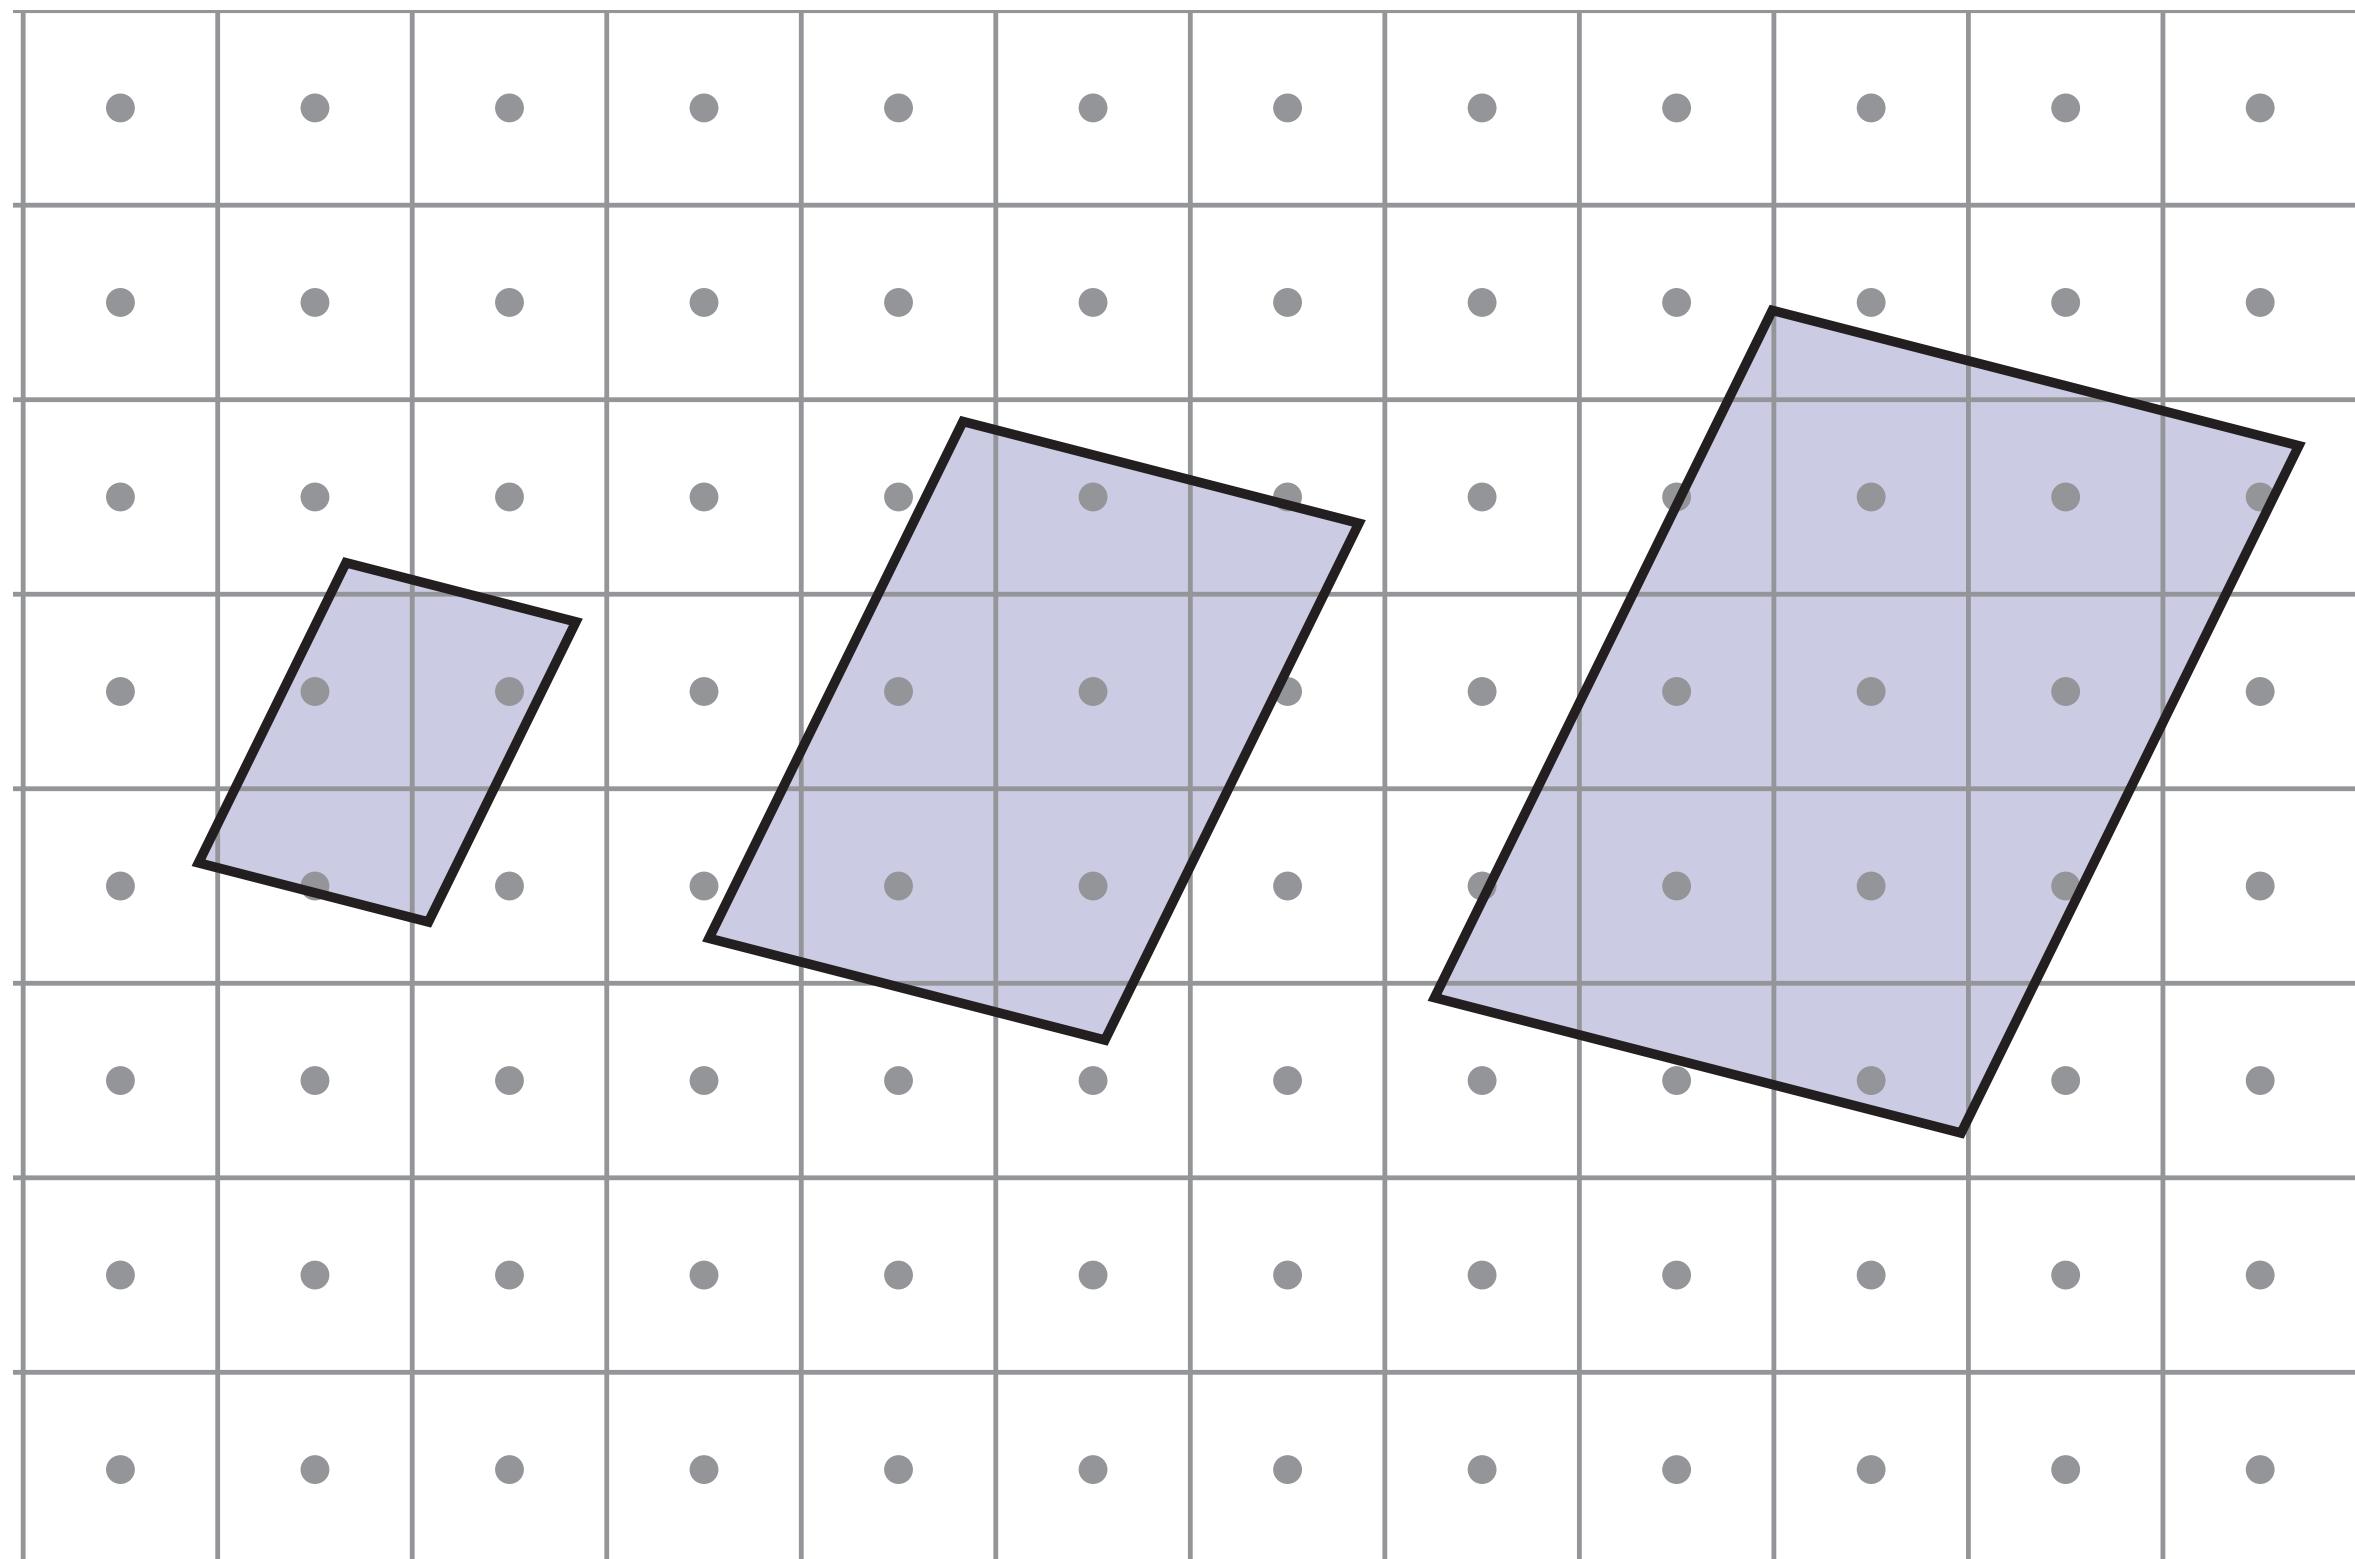
- Amount of pre-filtering depends on how far away the object is:
 - minor minification: image pixel extreme magnification: image pixel spans large region of texture
- Idea:
 - Low-pass filter and downsample texture file, and store successively lower resolutions
 - For each sample, use the texture file whose resolution approximates the screen sampling rate



Shader region = pixel area
Red lines = screen pixel boundaries
Red dots = texture space sample points for adjacent pixels

But how much should we pre-filter?

- **Amount of pre-filtering necessary depends on how far away the object is**
- **Idea: pre-compute and store different versions of the texture with different amounts of prefiltering**
 - **Low-pass filter and downsample texture file, and store successively lower resolutions**
 - **When sampling texture, use the texture file whose prefiltering amount matches the desired sampling rate**

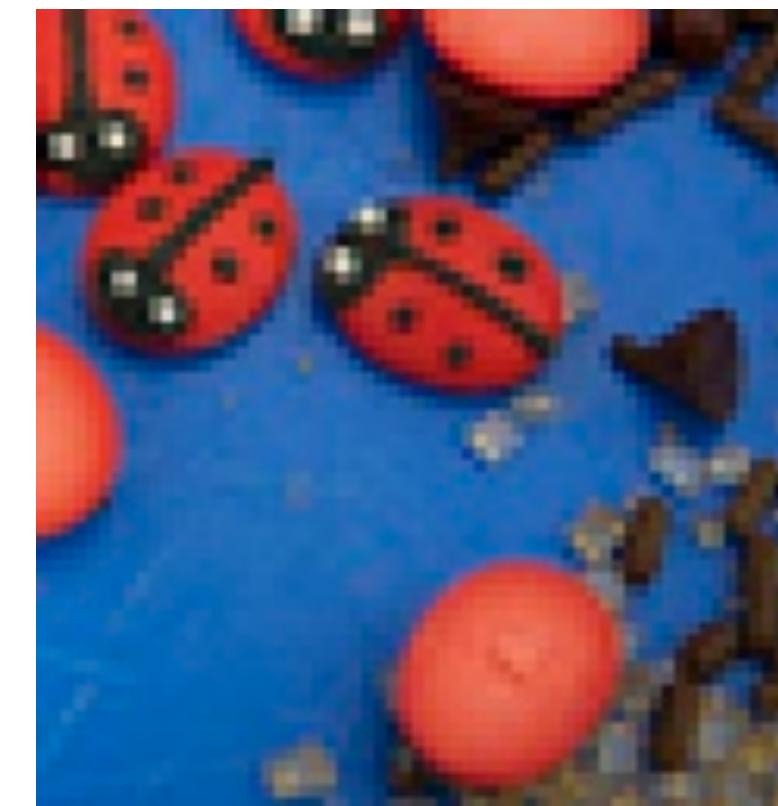


Mipmap (L. Williams 83)

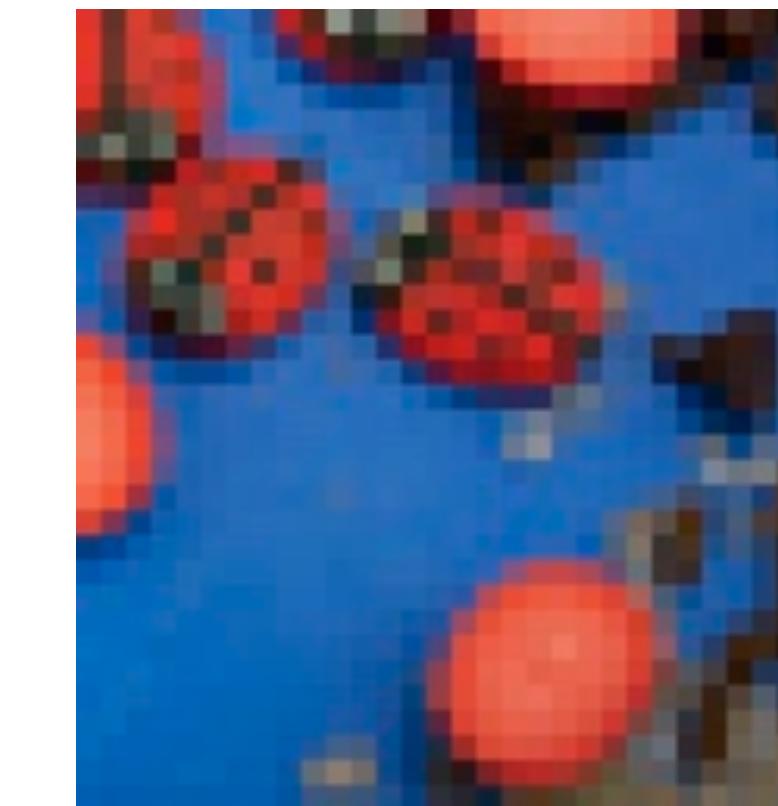
Each mipmap level is downsampled (low-pass filtered) version of the previous



Level 0 = 128x128



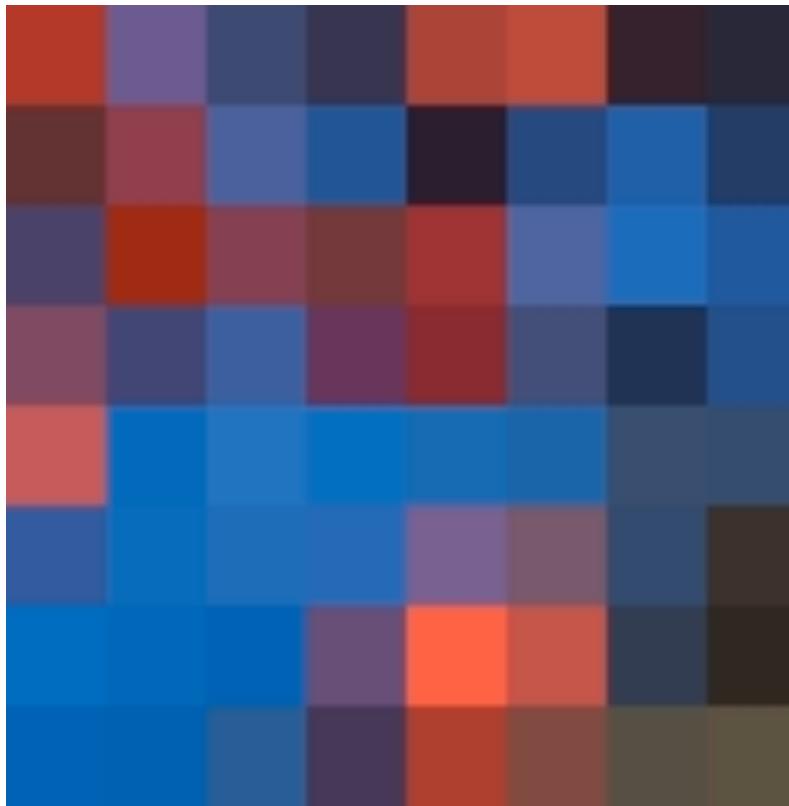
Level 1 = 64x64



Level 2 = 32x32



Level 3 = 16x16



Level 4 = 8x8



Level 5 = 4x4



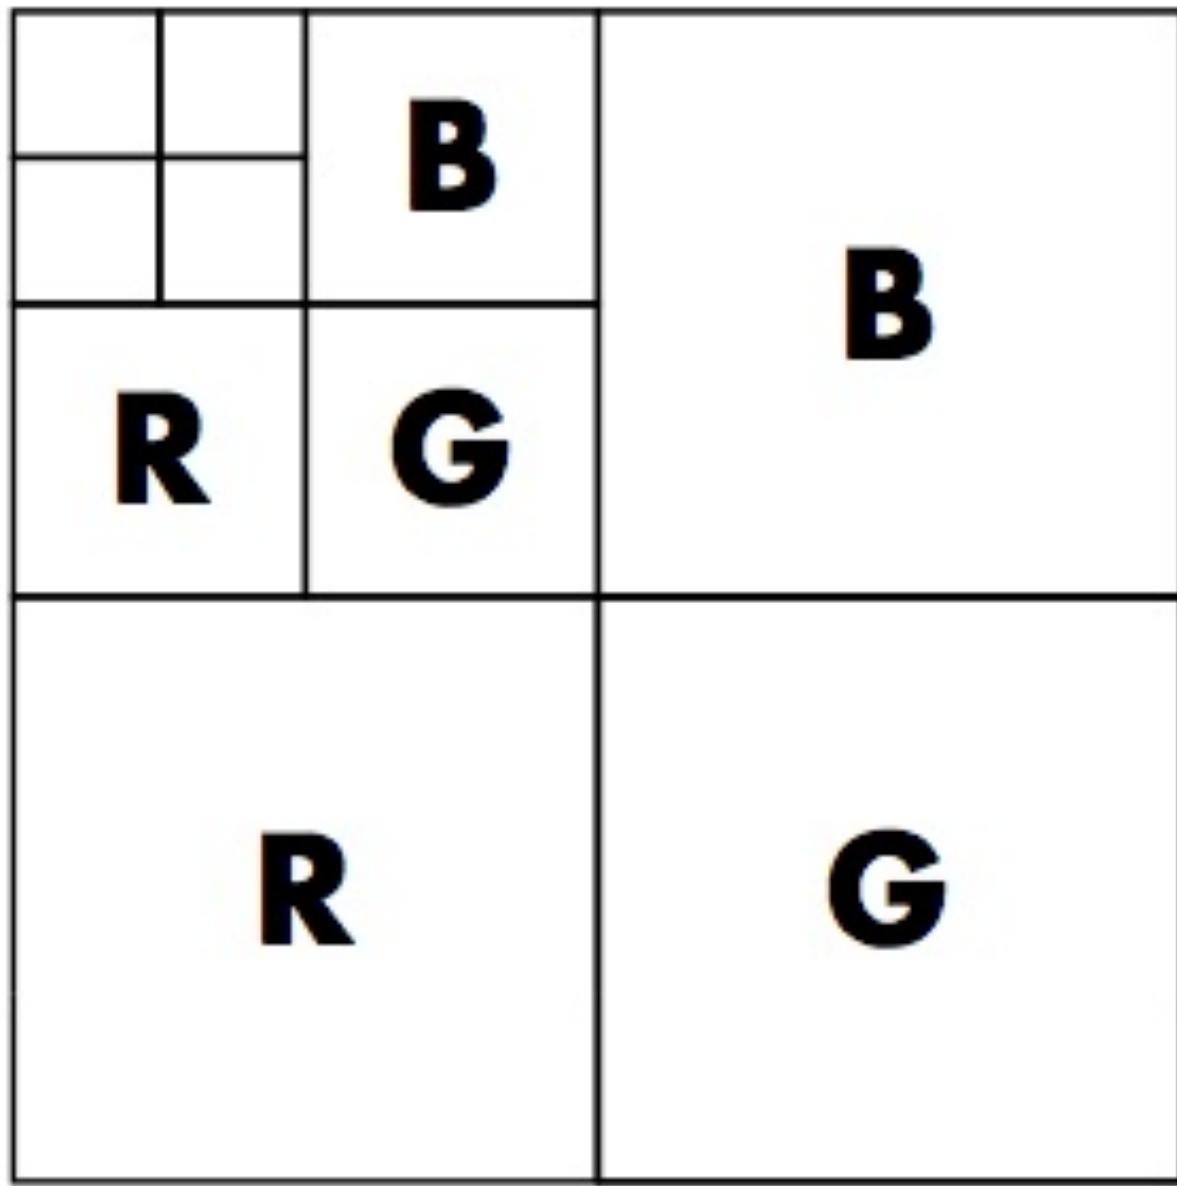
Level 6 = 2x2



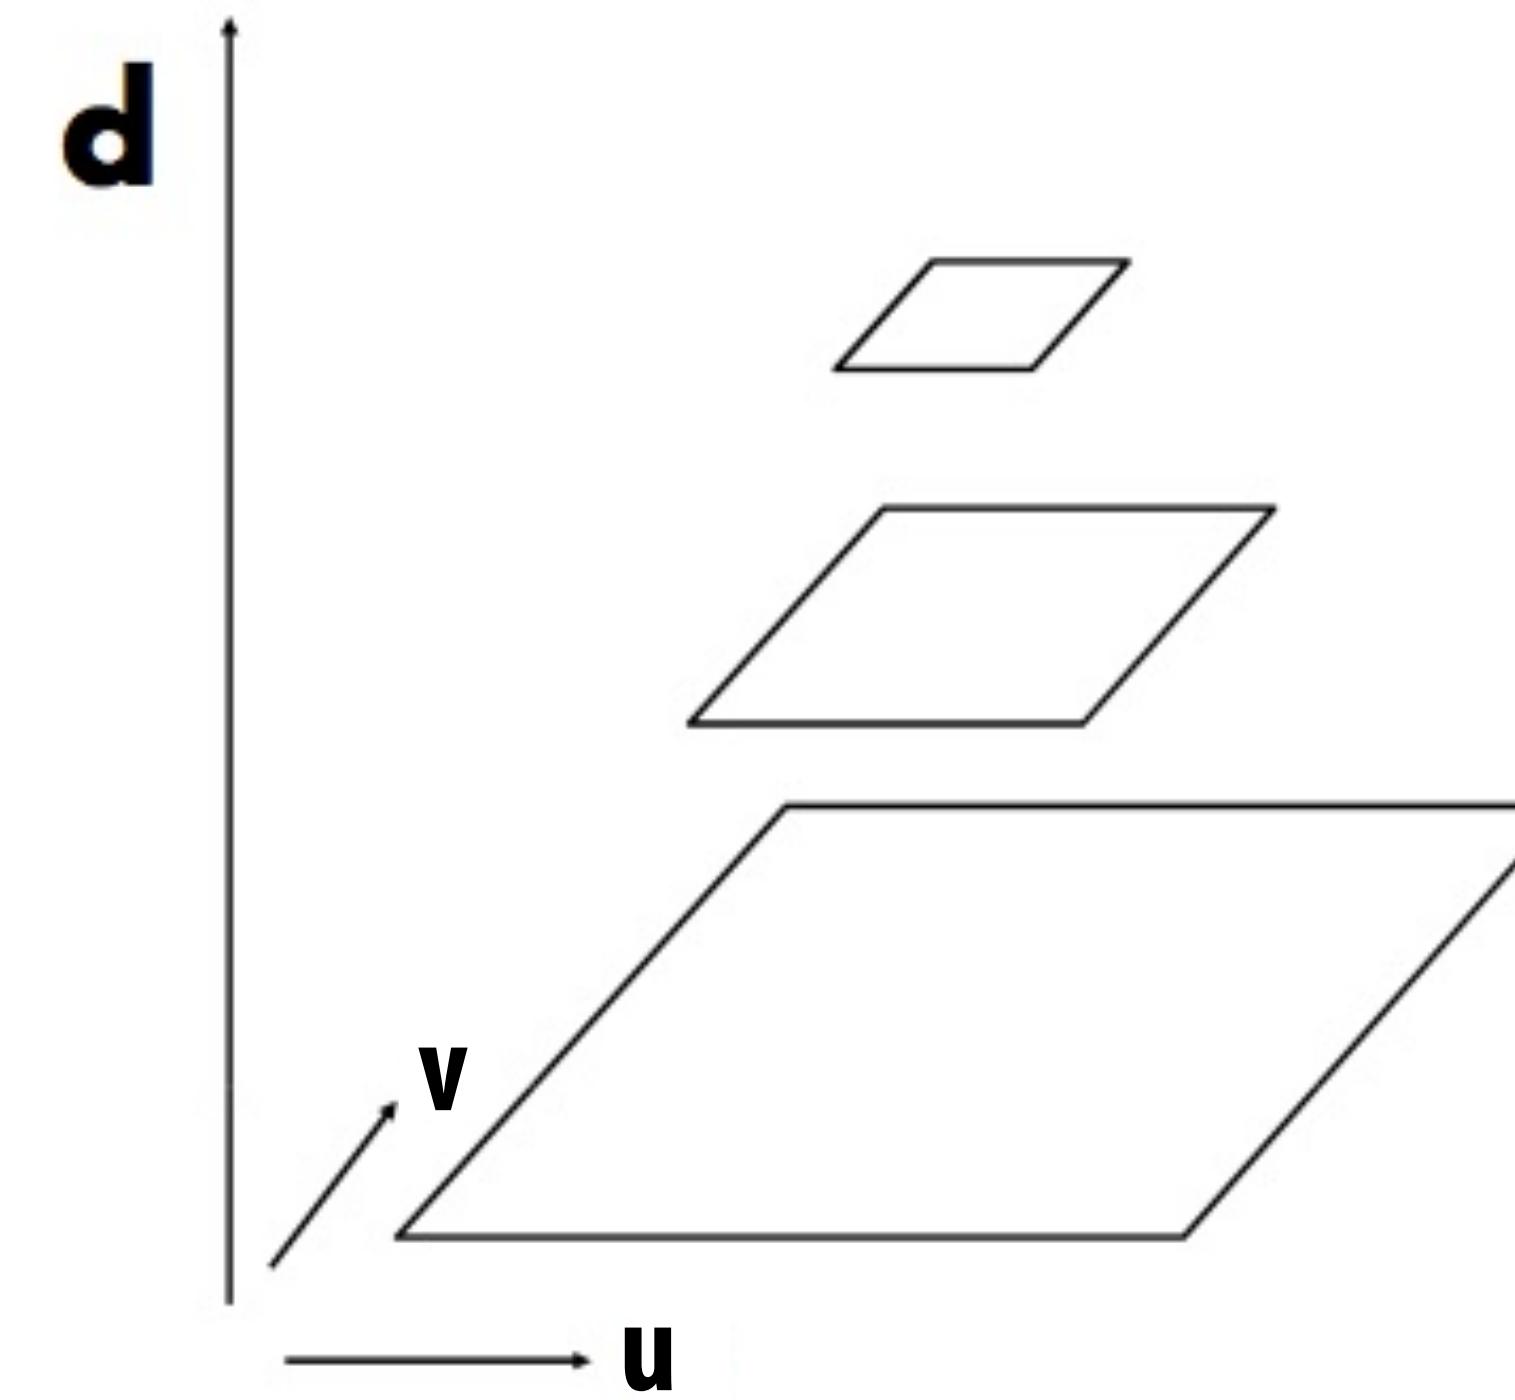
Level 7 = 1x1

“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

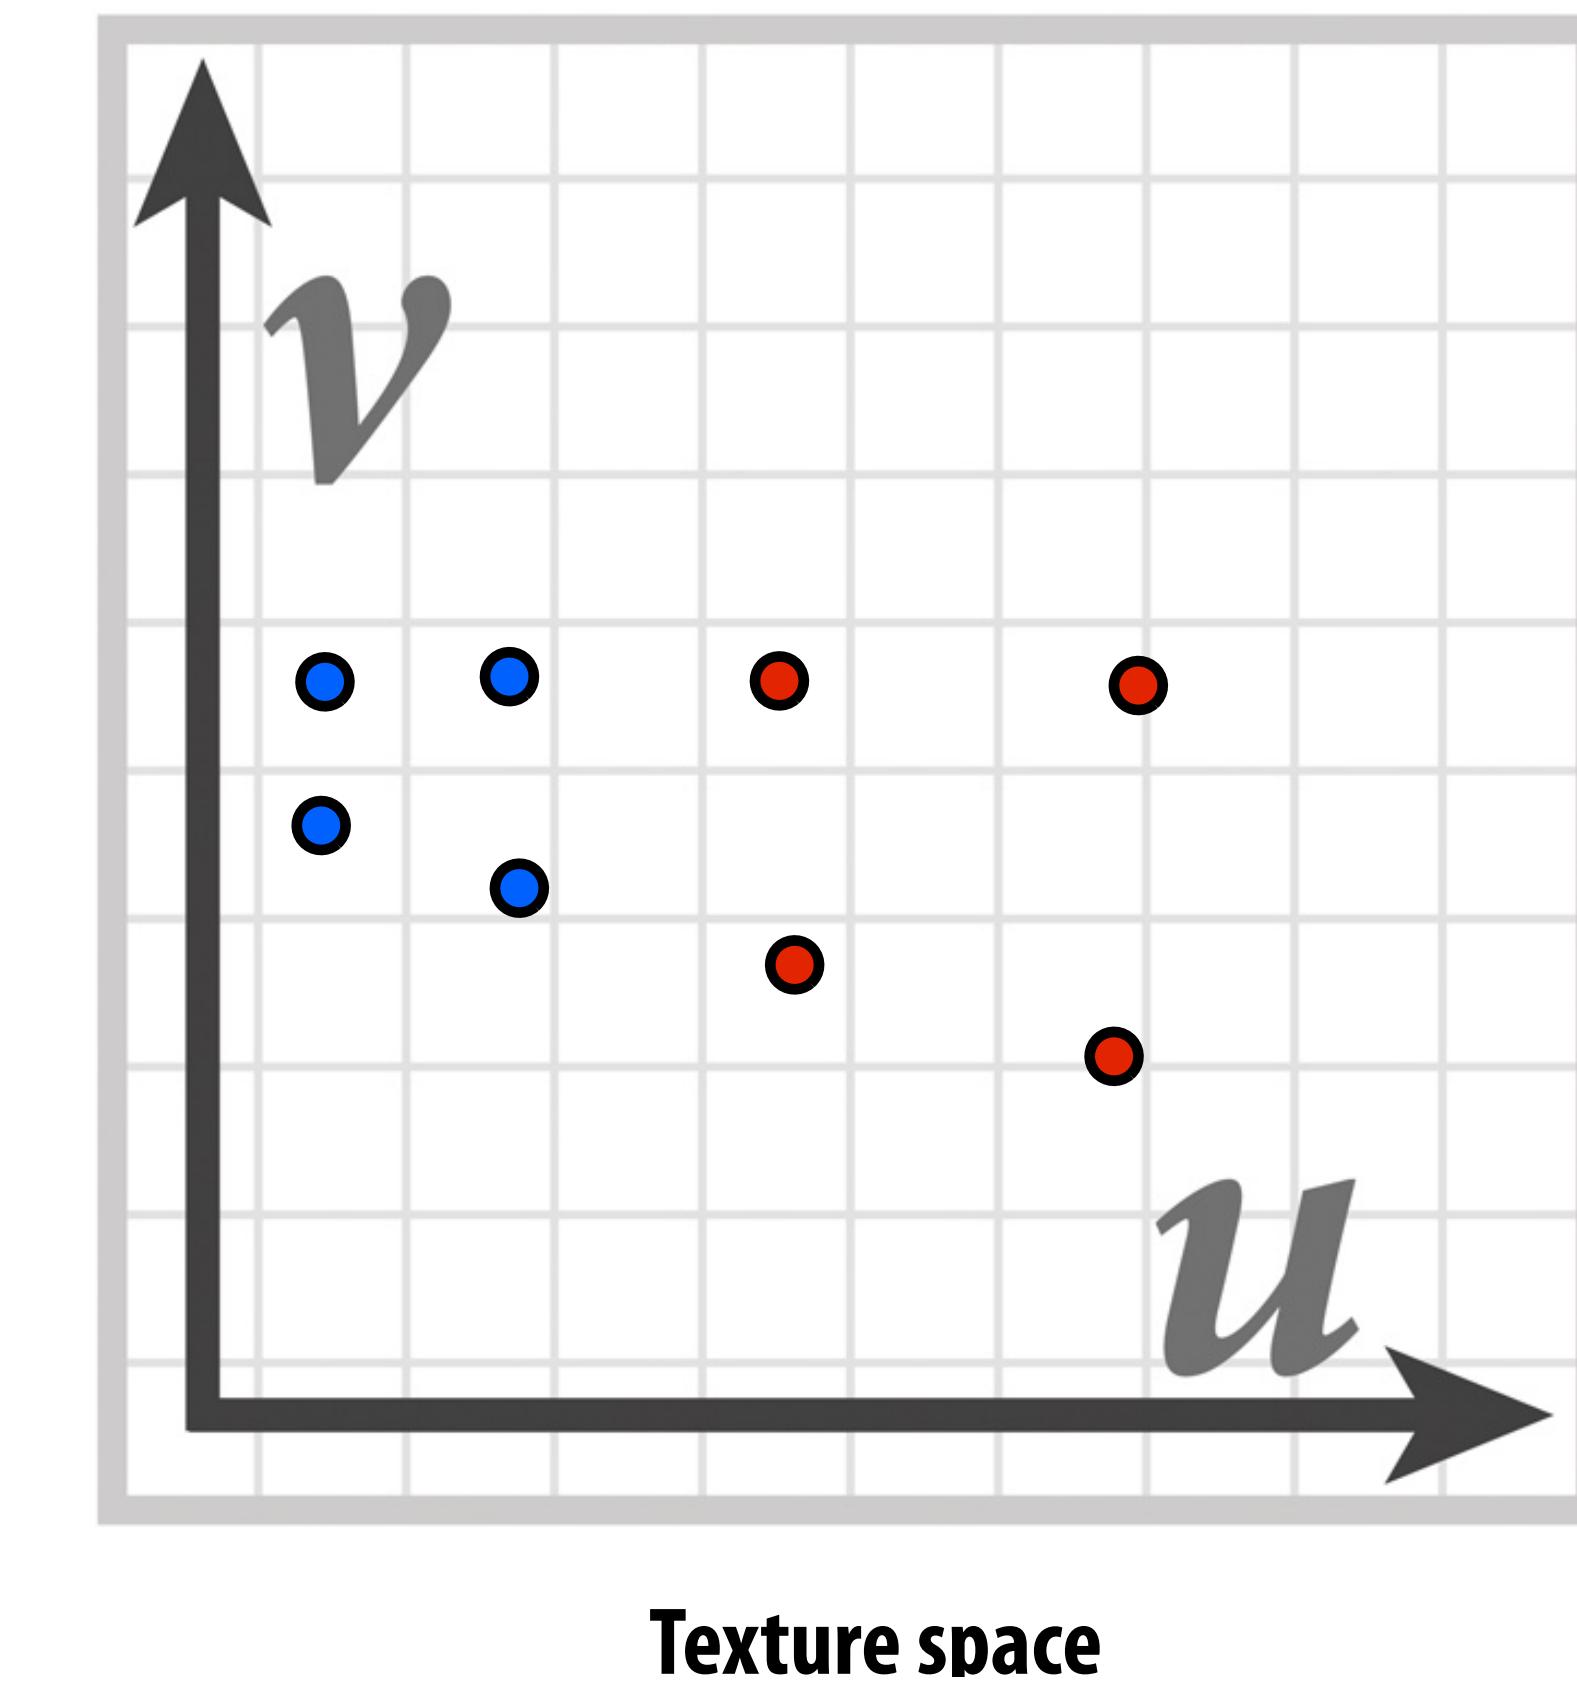
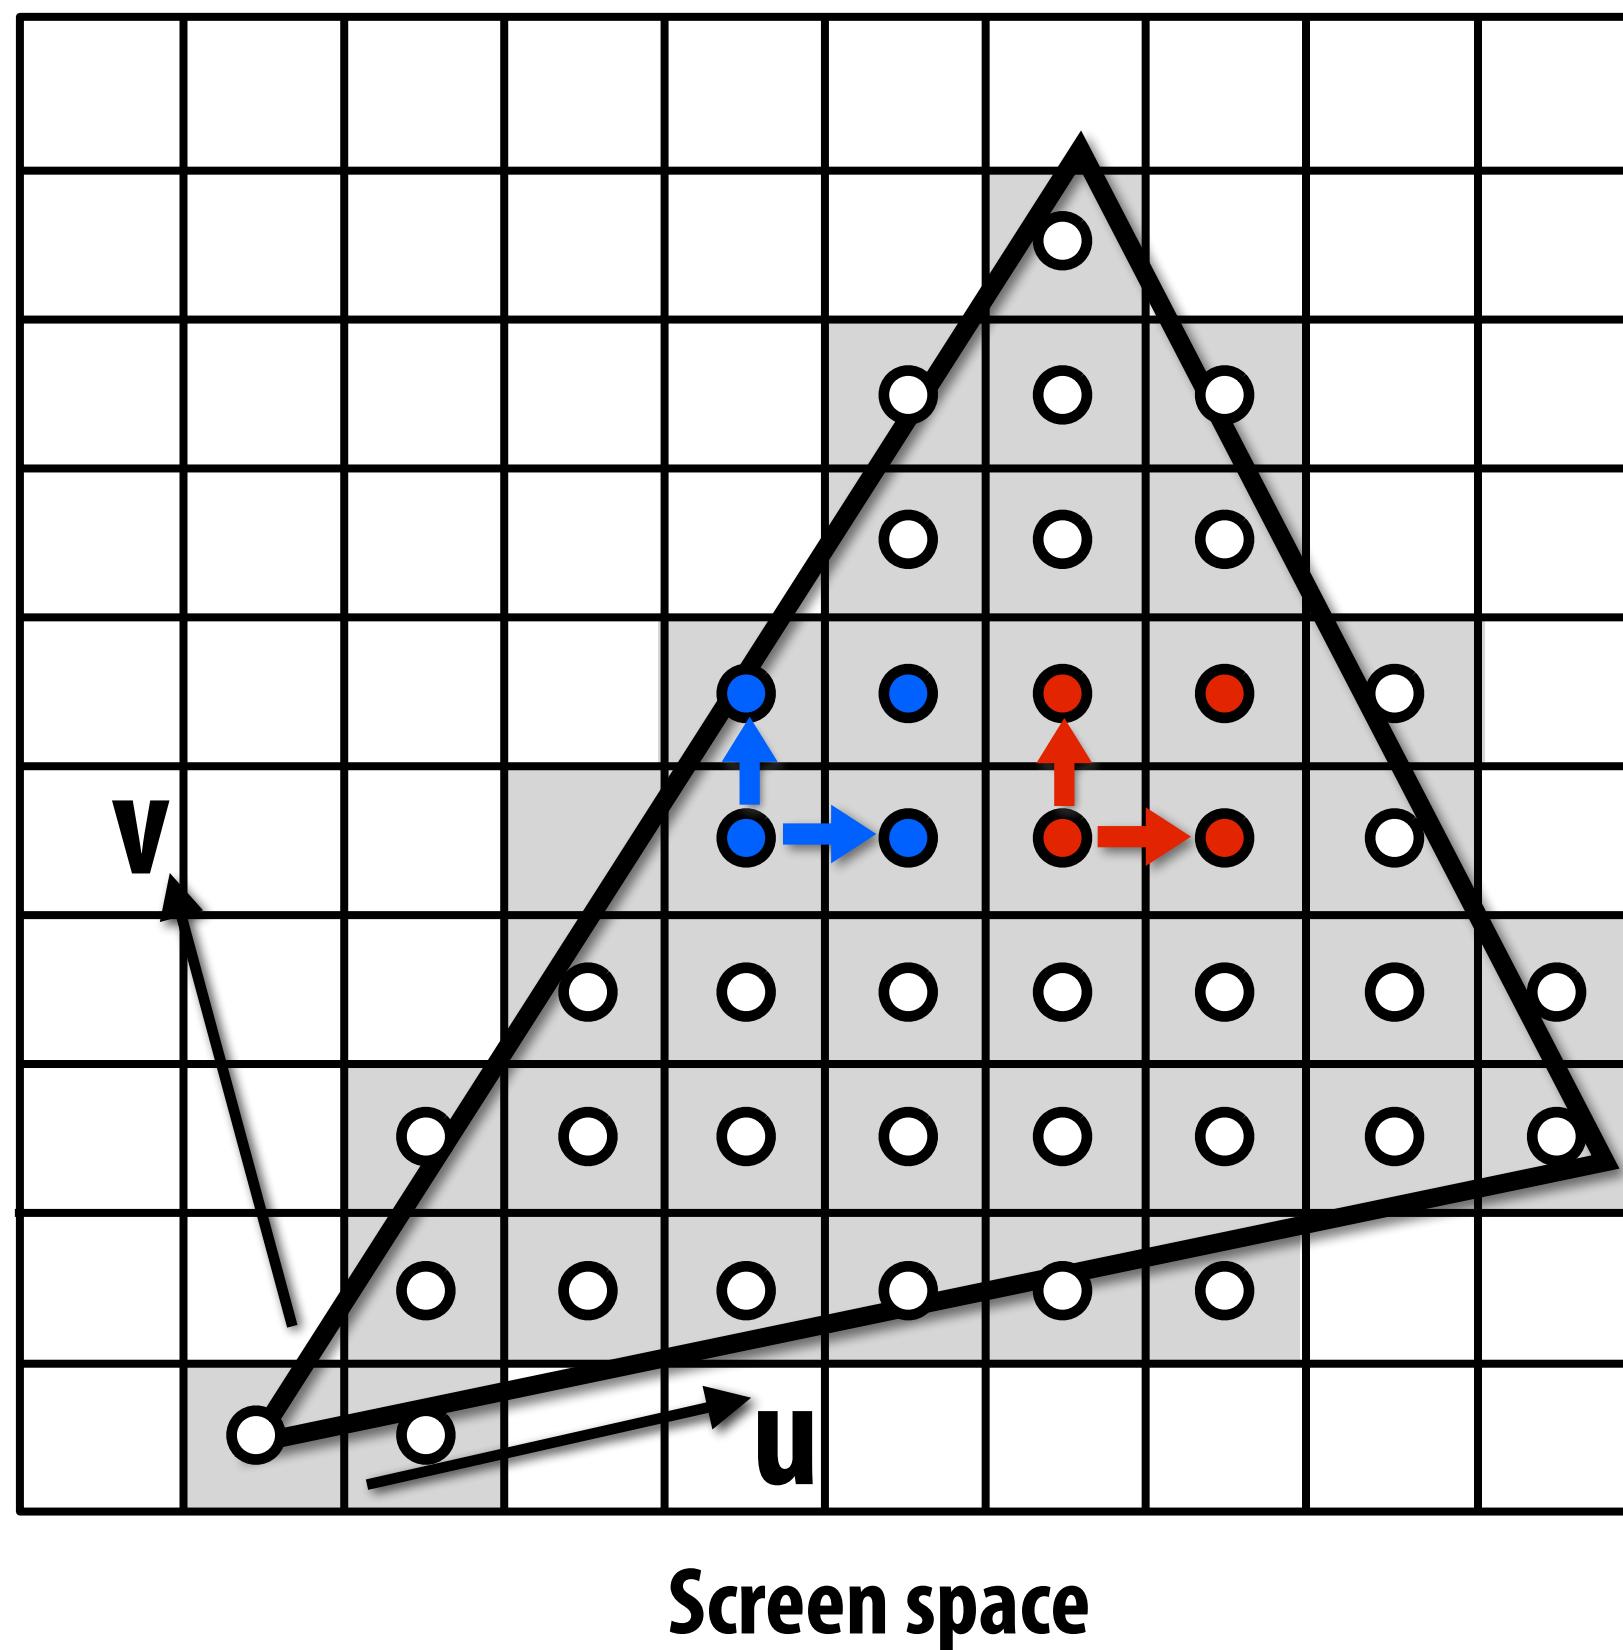


“Mip hierarchy”
level = d

What is the storage overhead of a mipmap?

Computing mipmap level (in a rasterizer) *

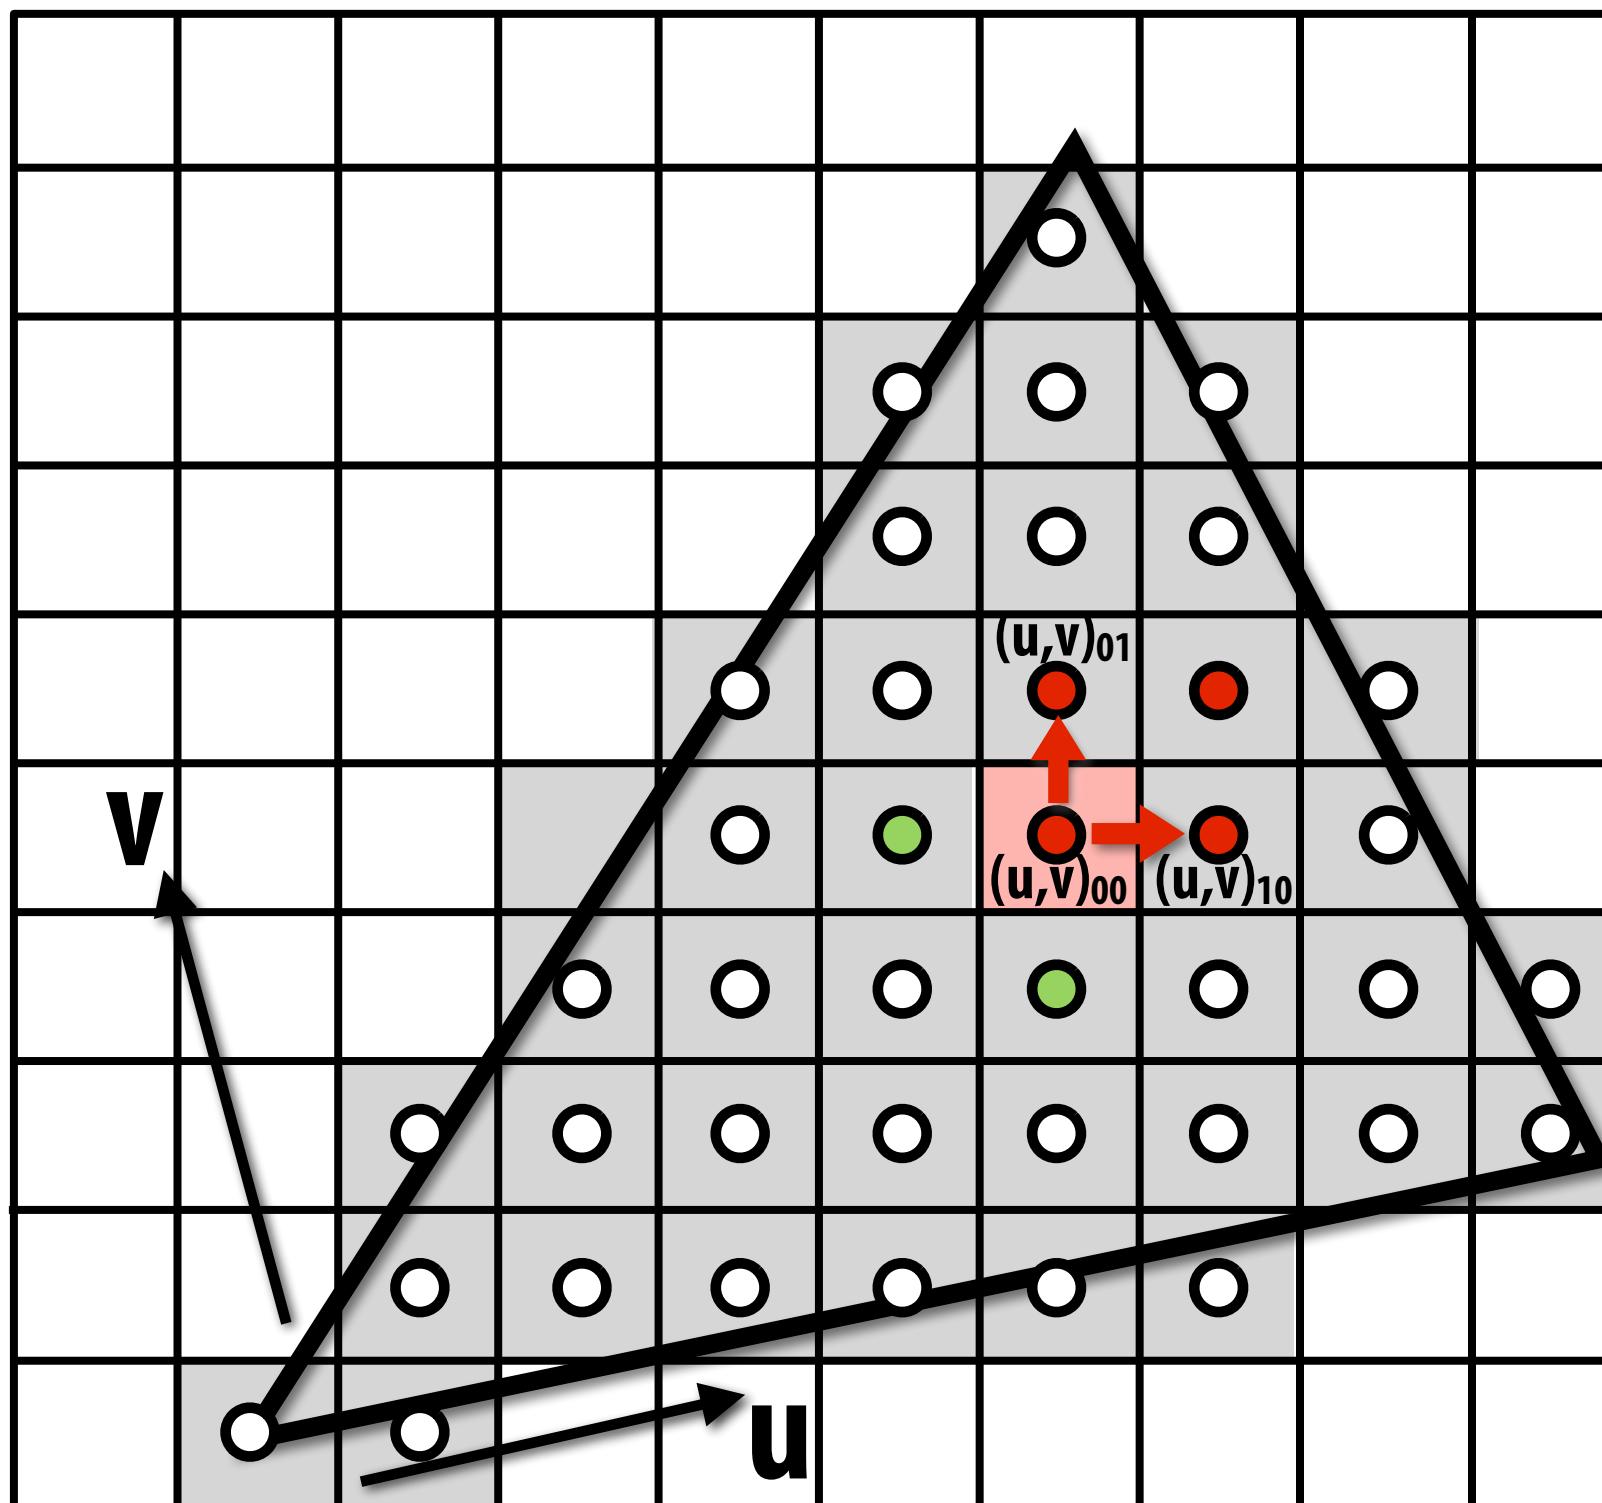
Compute differences between texture coordinate values of neighboring screen samples



* In assignment 2, we'll teach you about how this is done with ray differentials in a ray tracer

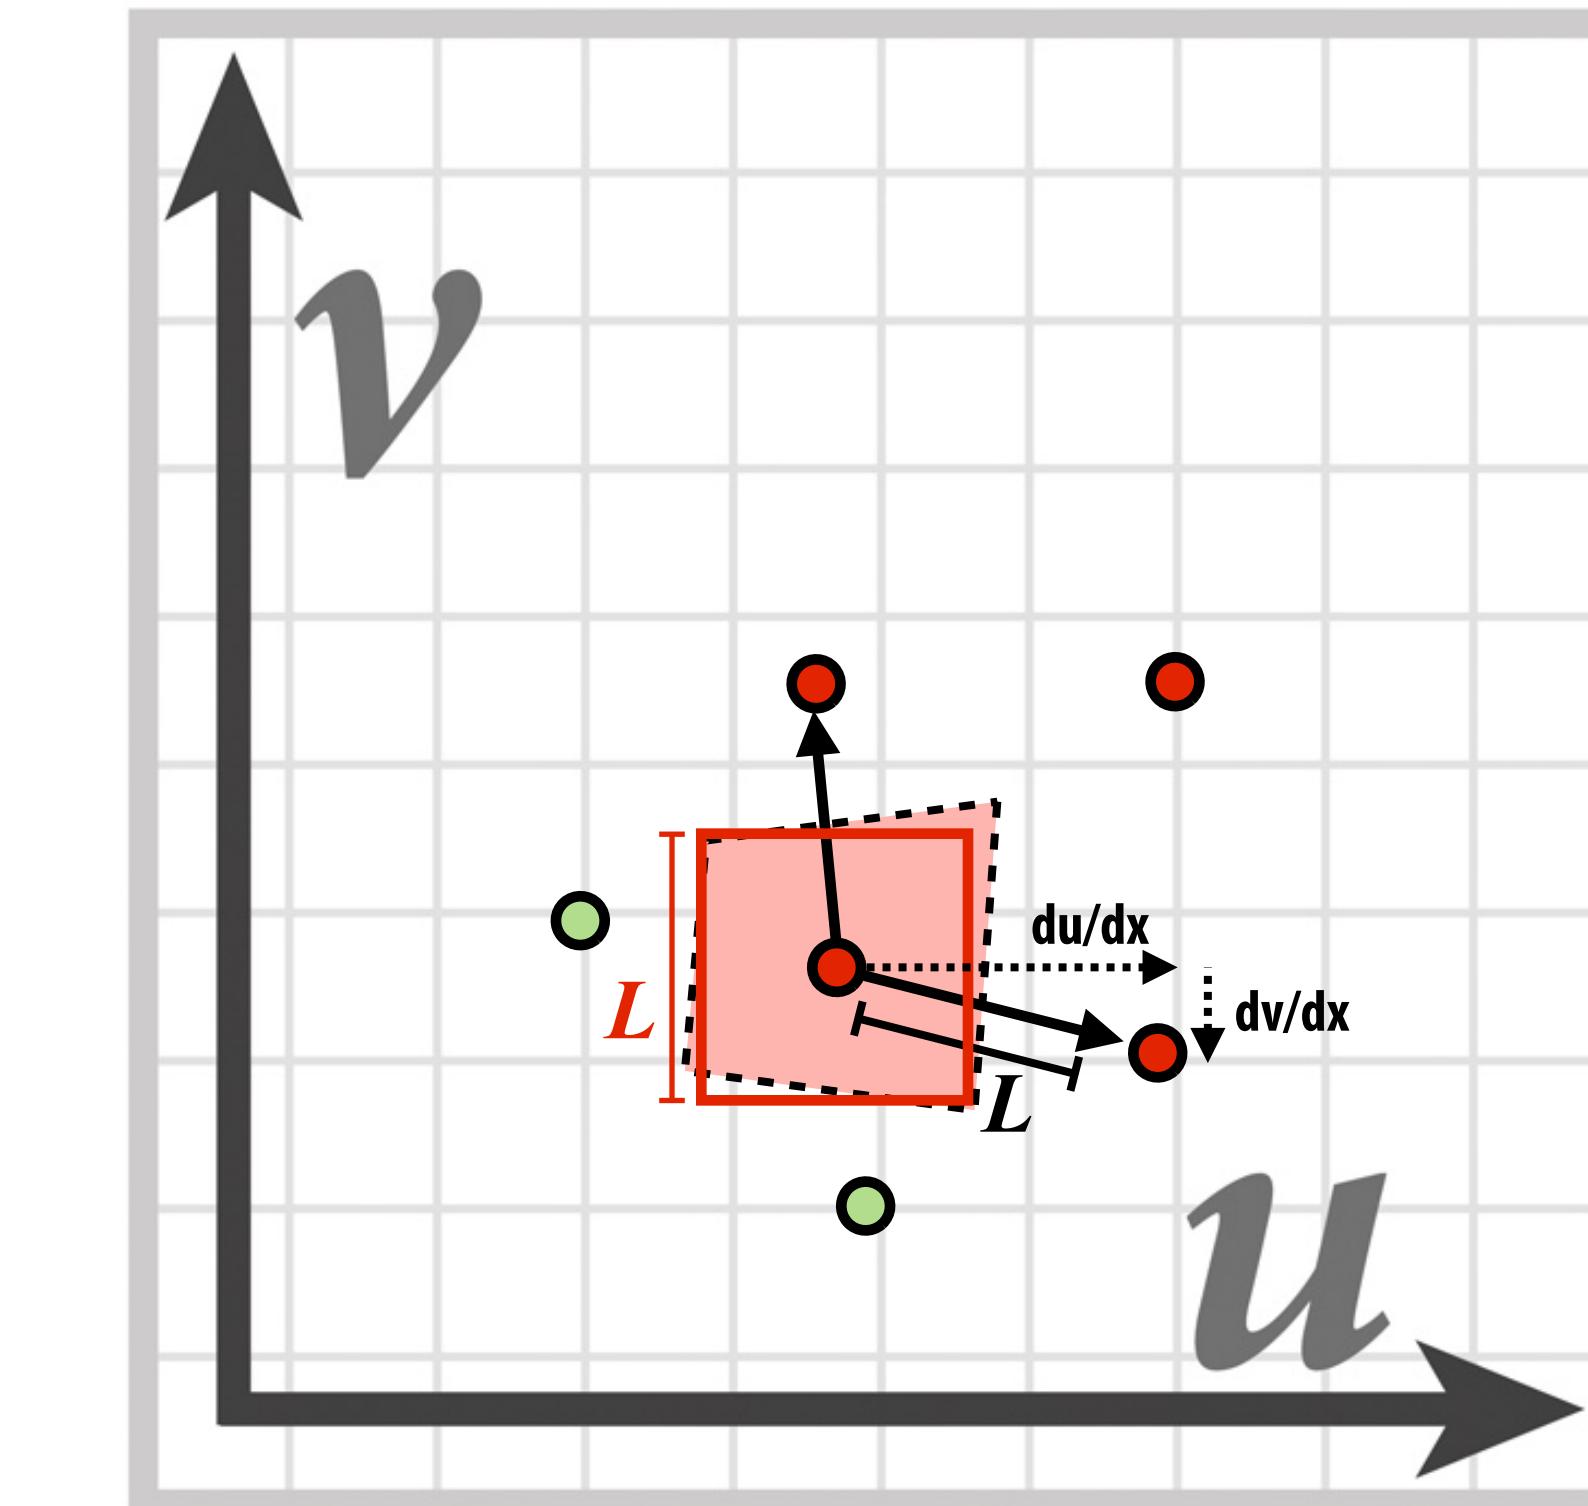
Computing mipmap level (in a rasterizer) *

Compute differences between texture coordinate values of neighboring screen samples



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

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



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

mip-map d = $\log_2 L$

* In assignment 2, we'll teach you about how this is done with ray differentials in a ray tracer

Bilinear resampling at level 0



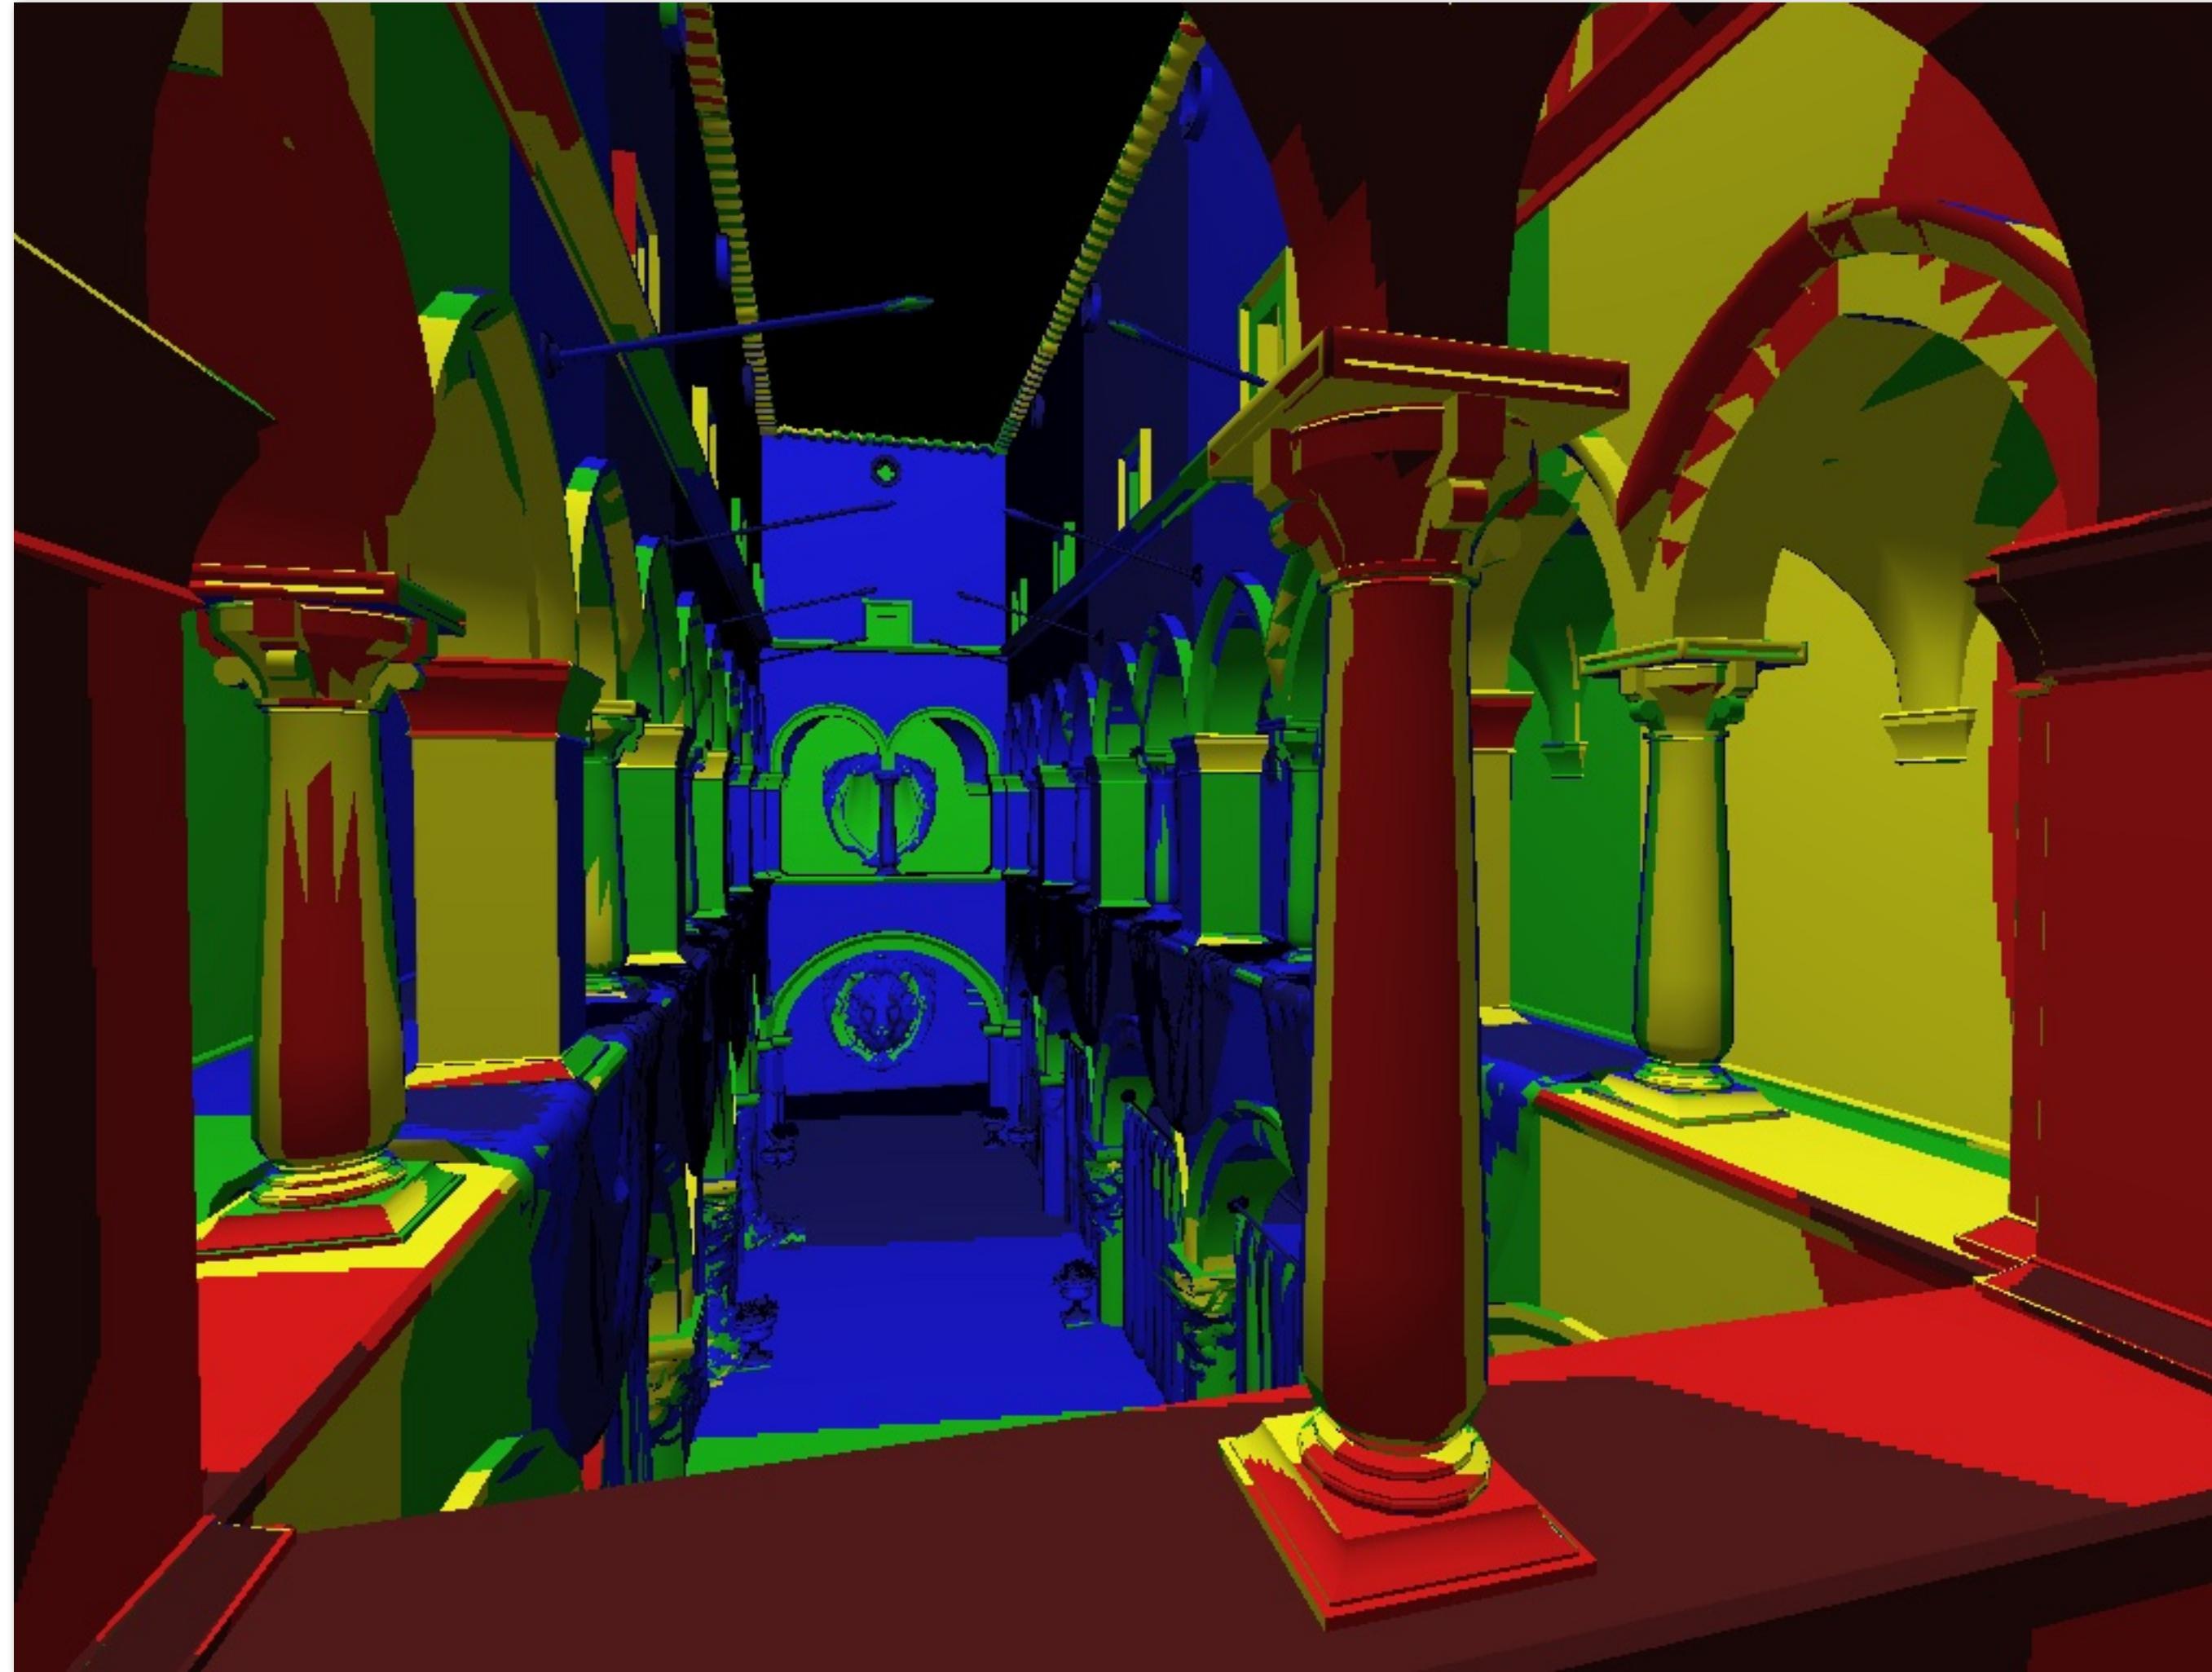
Bilinear resampling at level 2



Bilinear resampling at level 4



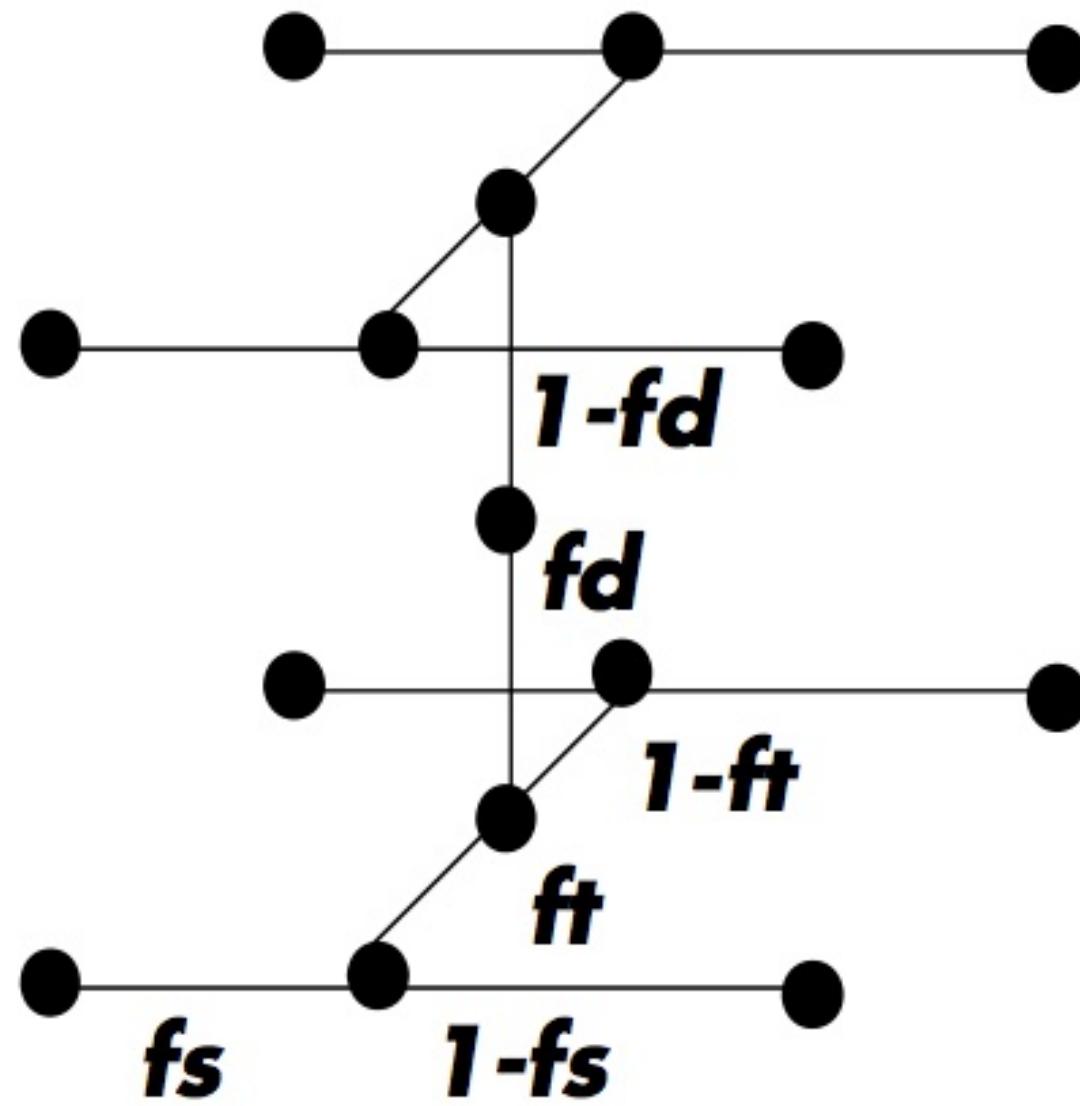
Visualization of mipmap level (bilinear filtering only: d clamped to nearest level)



“Tri-linear” filtering

Linearly interpolate the bilinear interpolation results from two adjacent levels of the mip map.

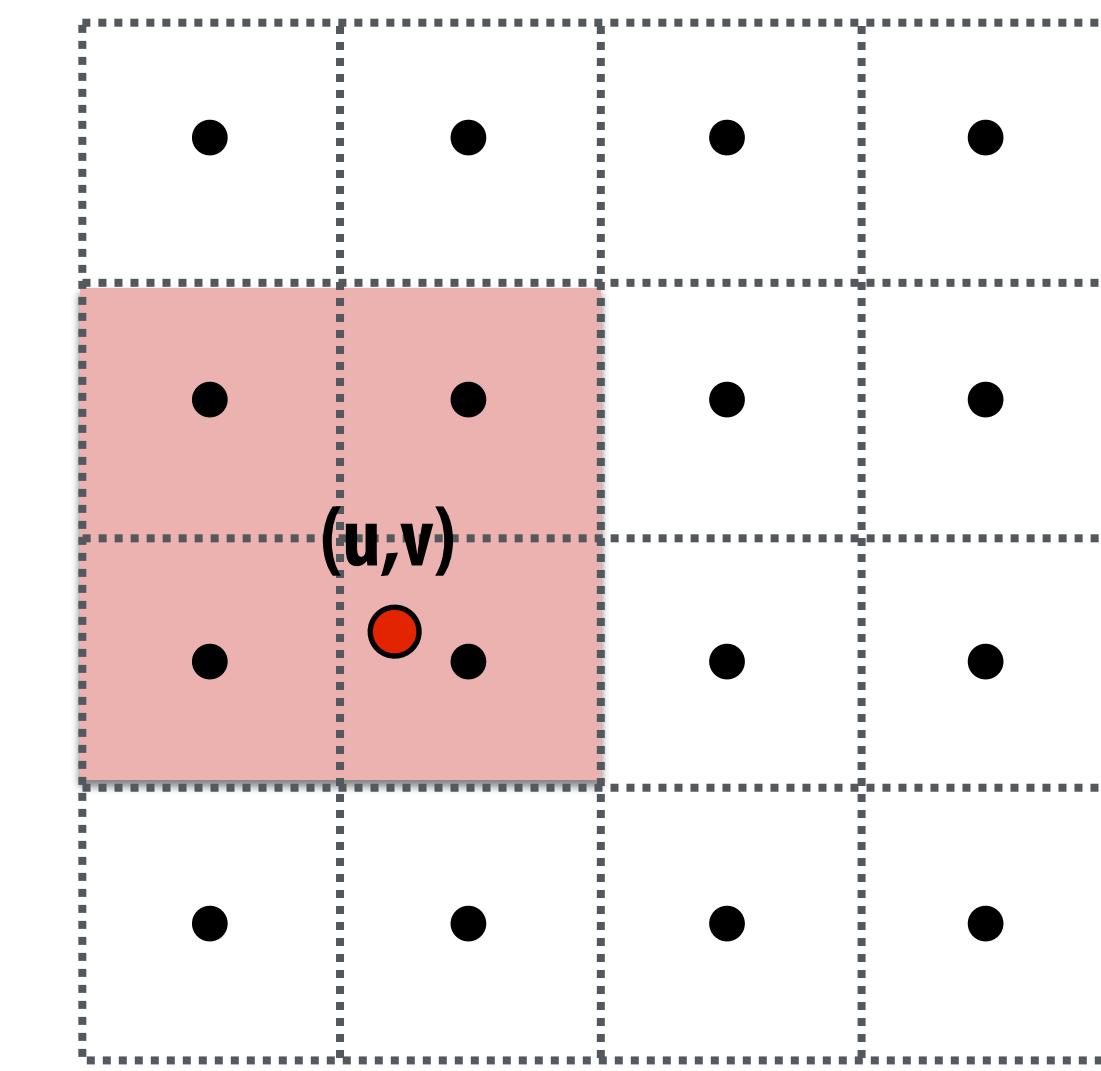
(smoothly transition between different levels of prefiltering)



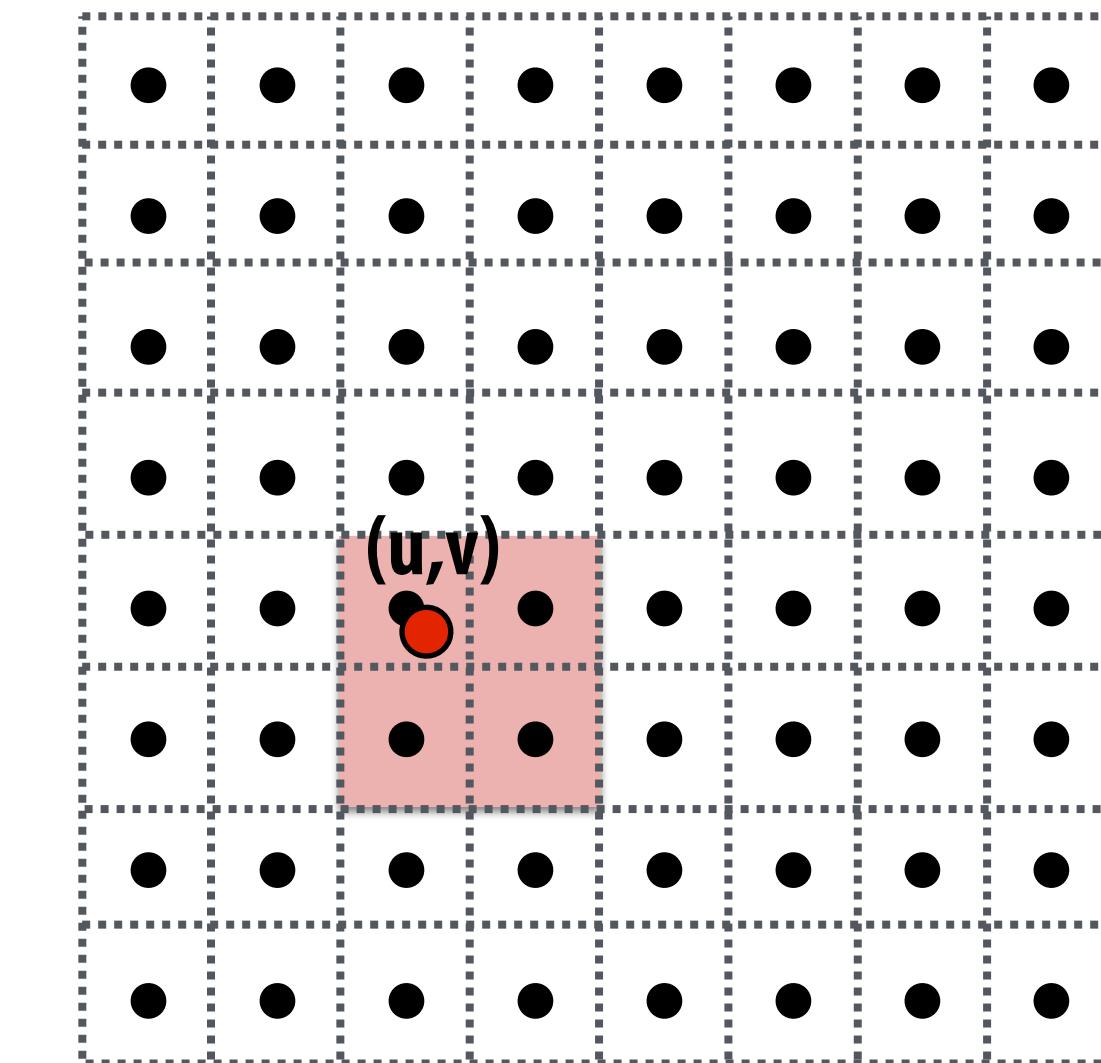
$$lerp(t, v_1, v_2) = v_1 + t(v_2 - v_1)$$

Bilinear resampling:
four texel reads
3 lerps (3 mul + 6 add)

Trilinear resampling:
eight texel reads
7 lerps (7 mul + 14 add)

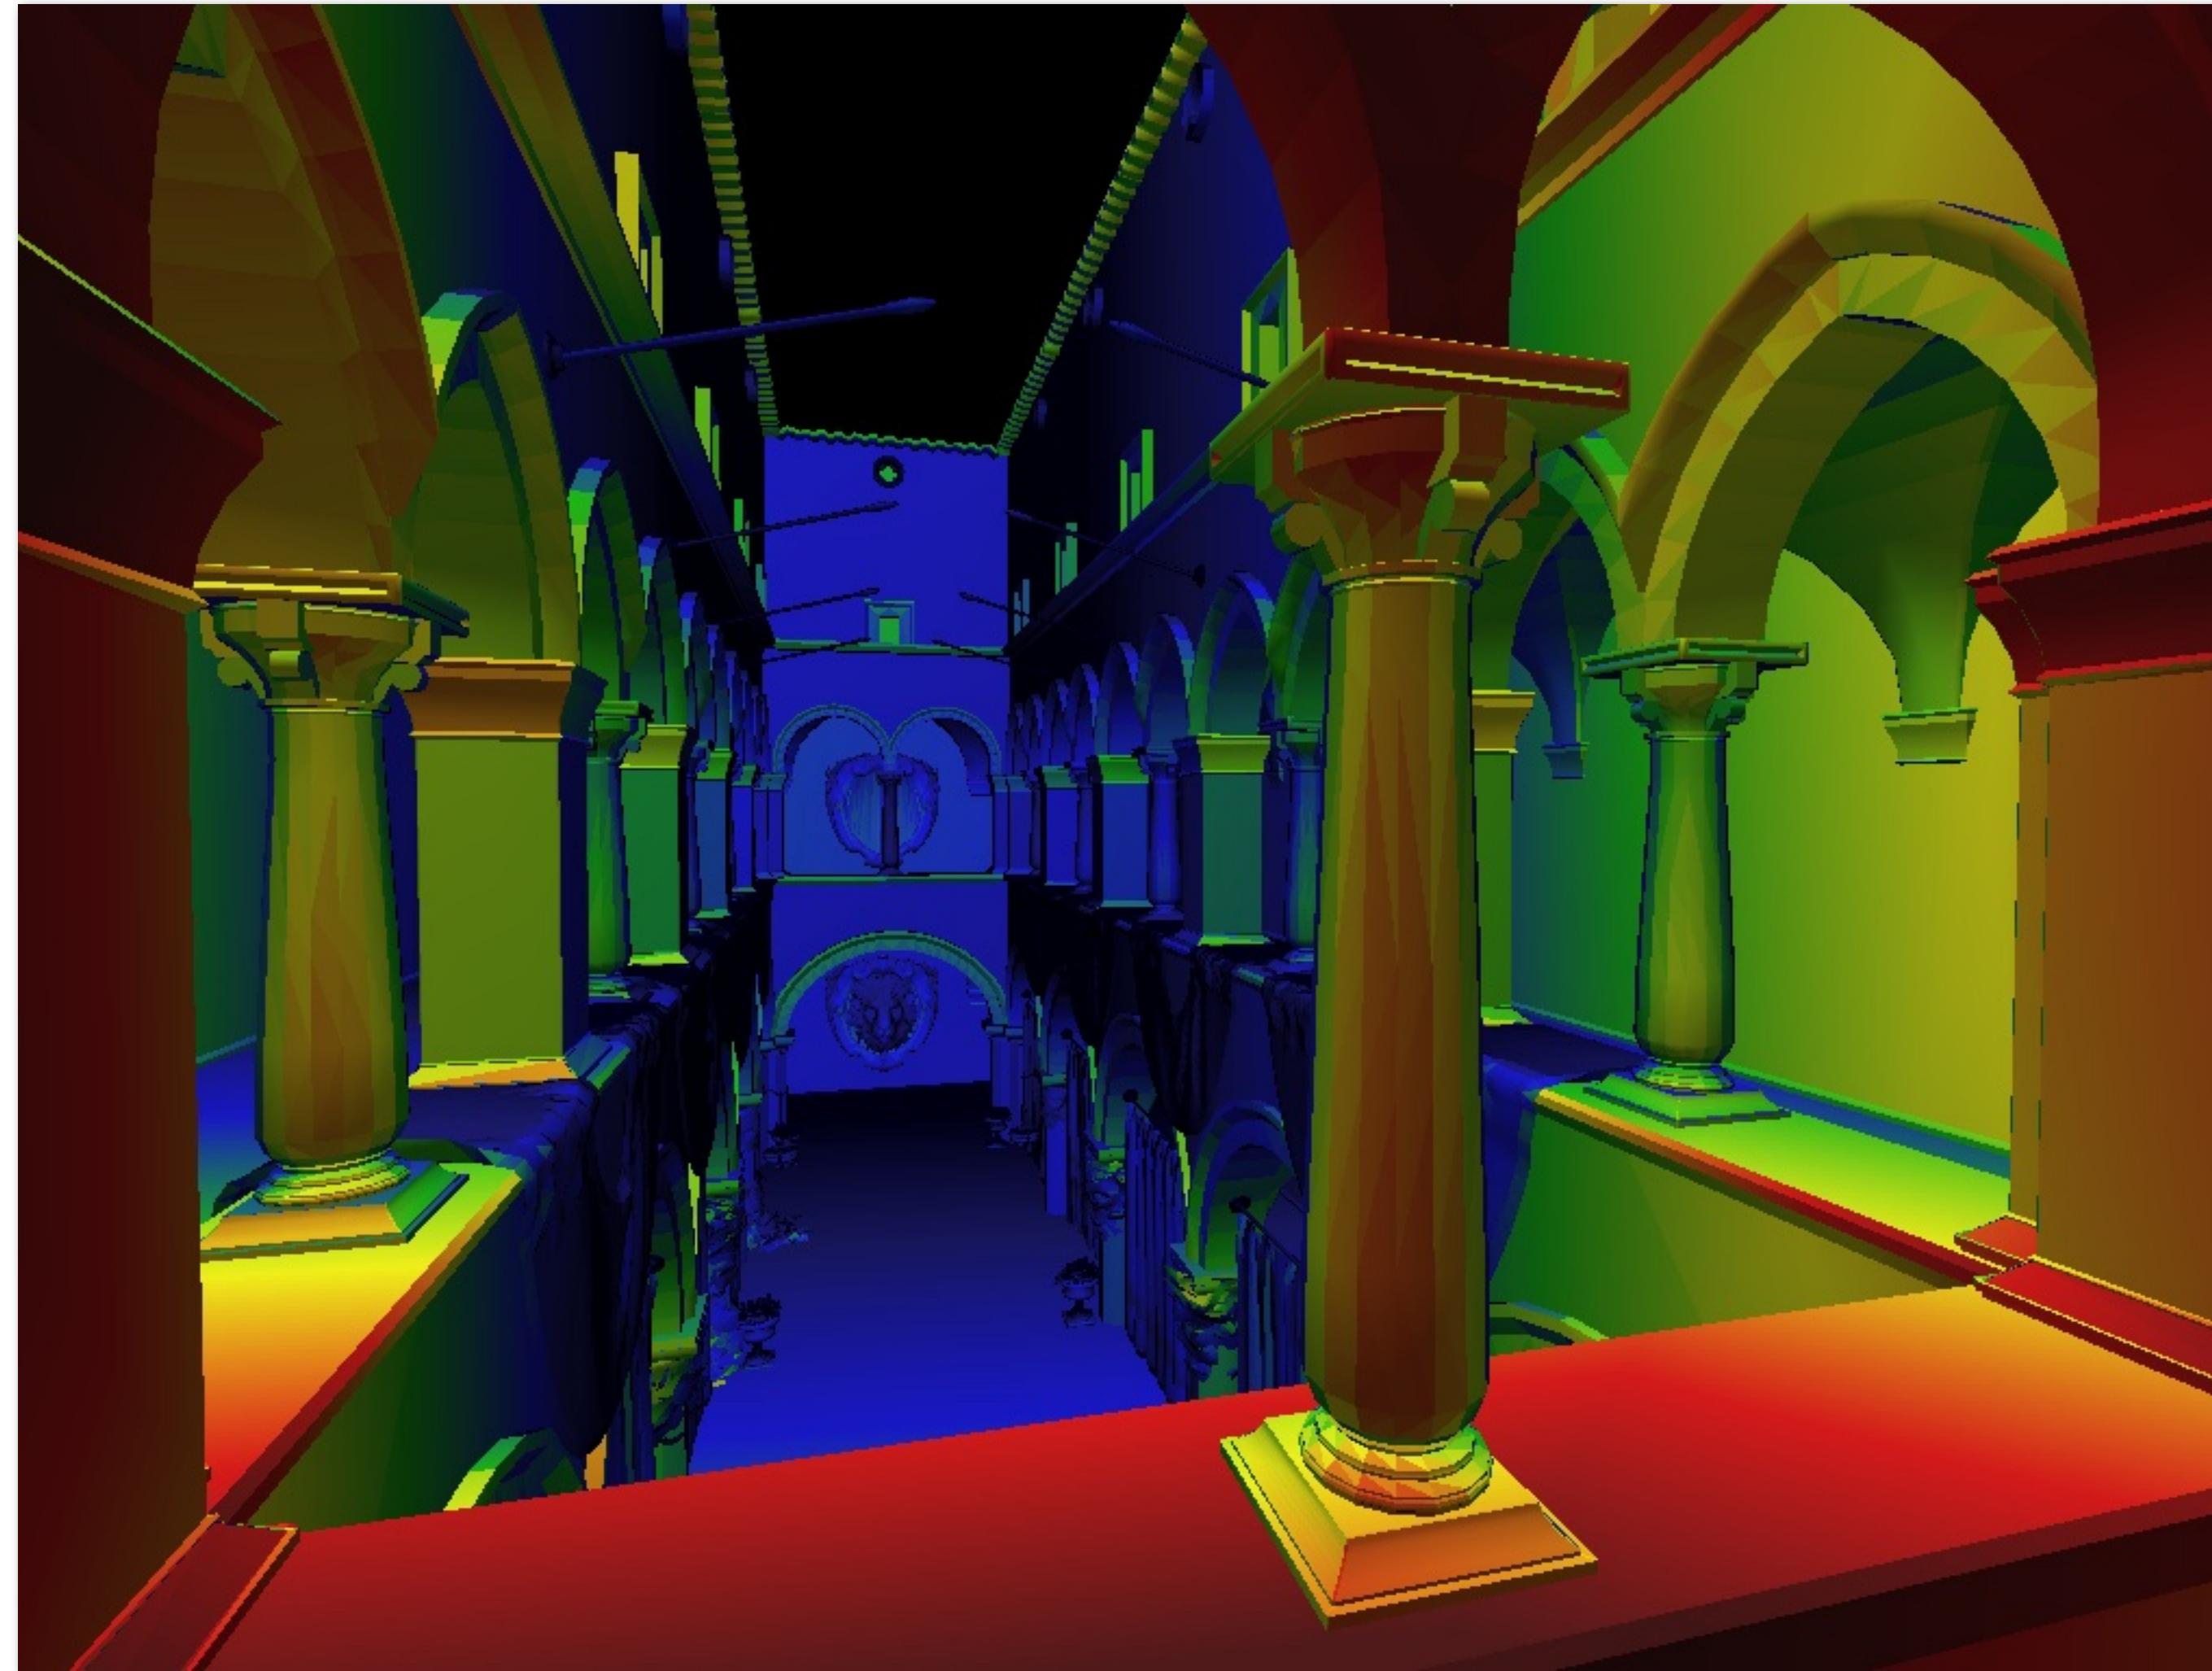


mip-map texels: level $d+1$



mip-map texels: level d

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



Bilinear vs trilinear filtering cost

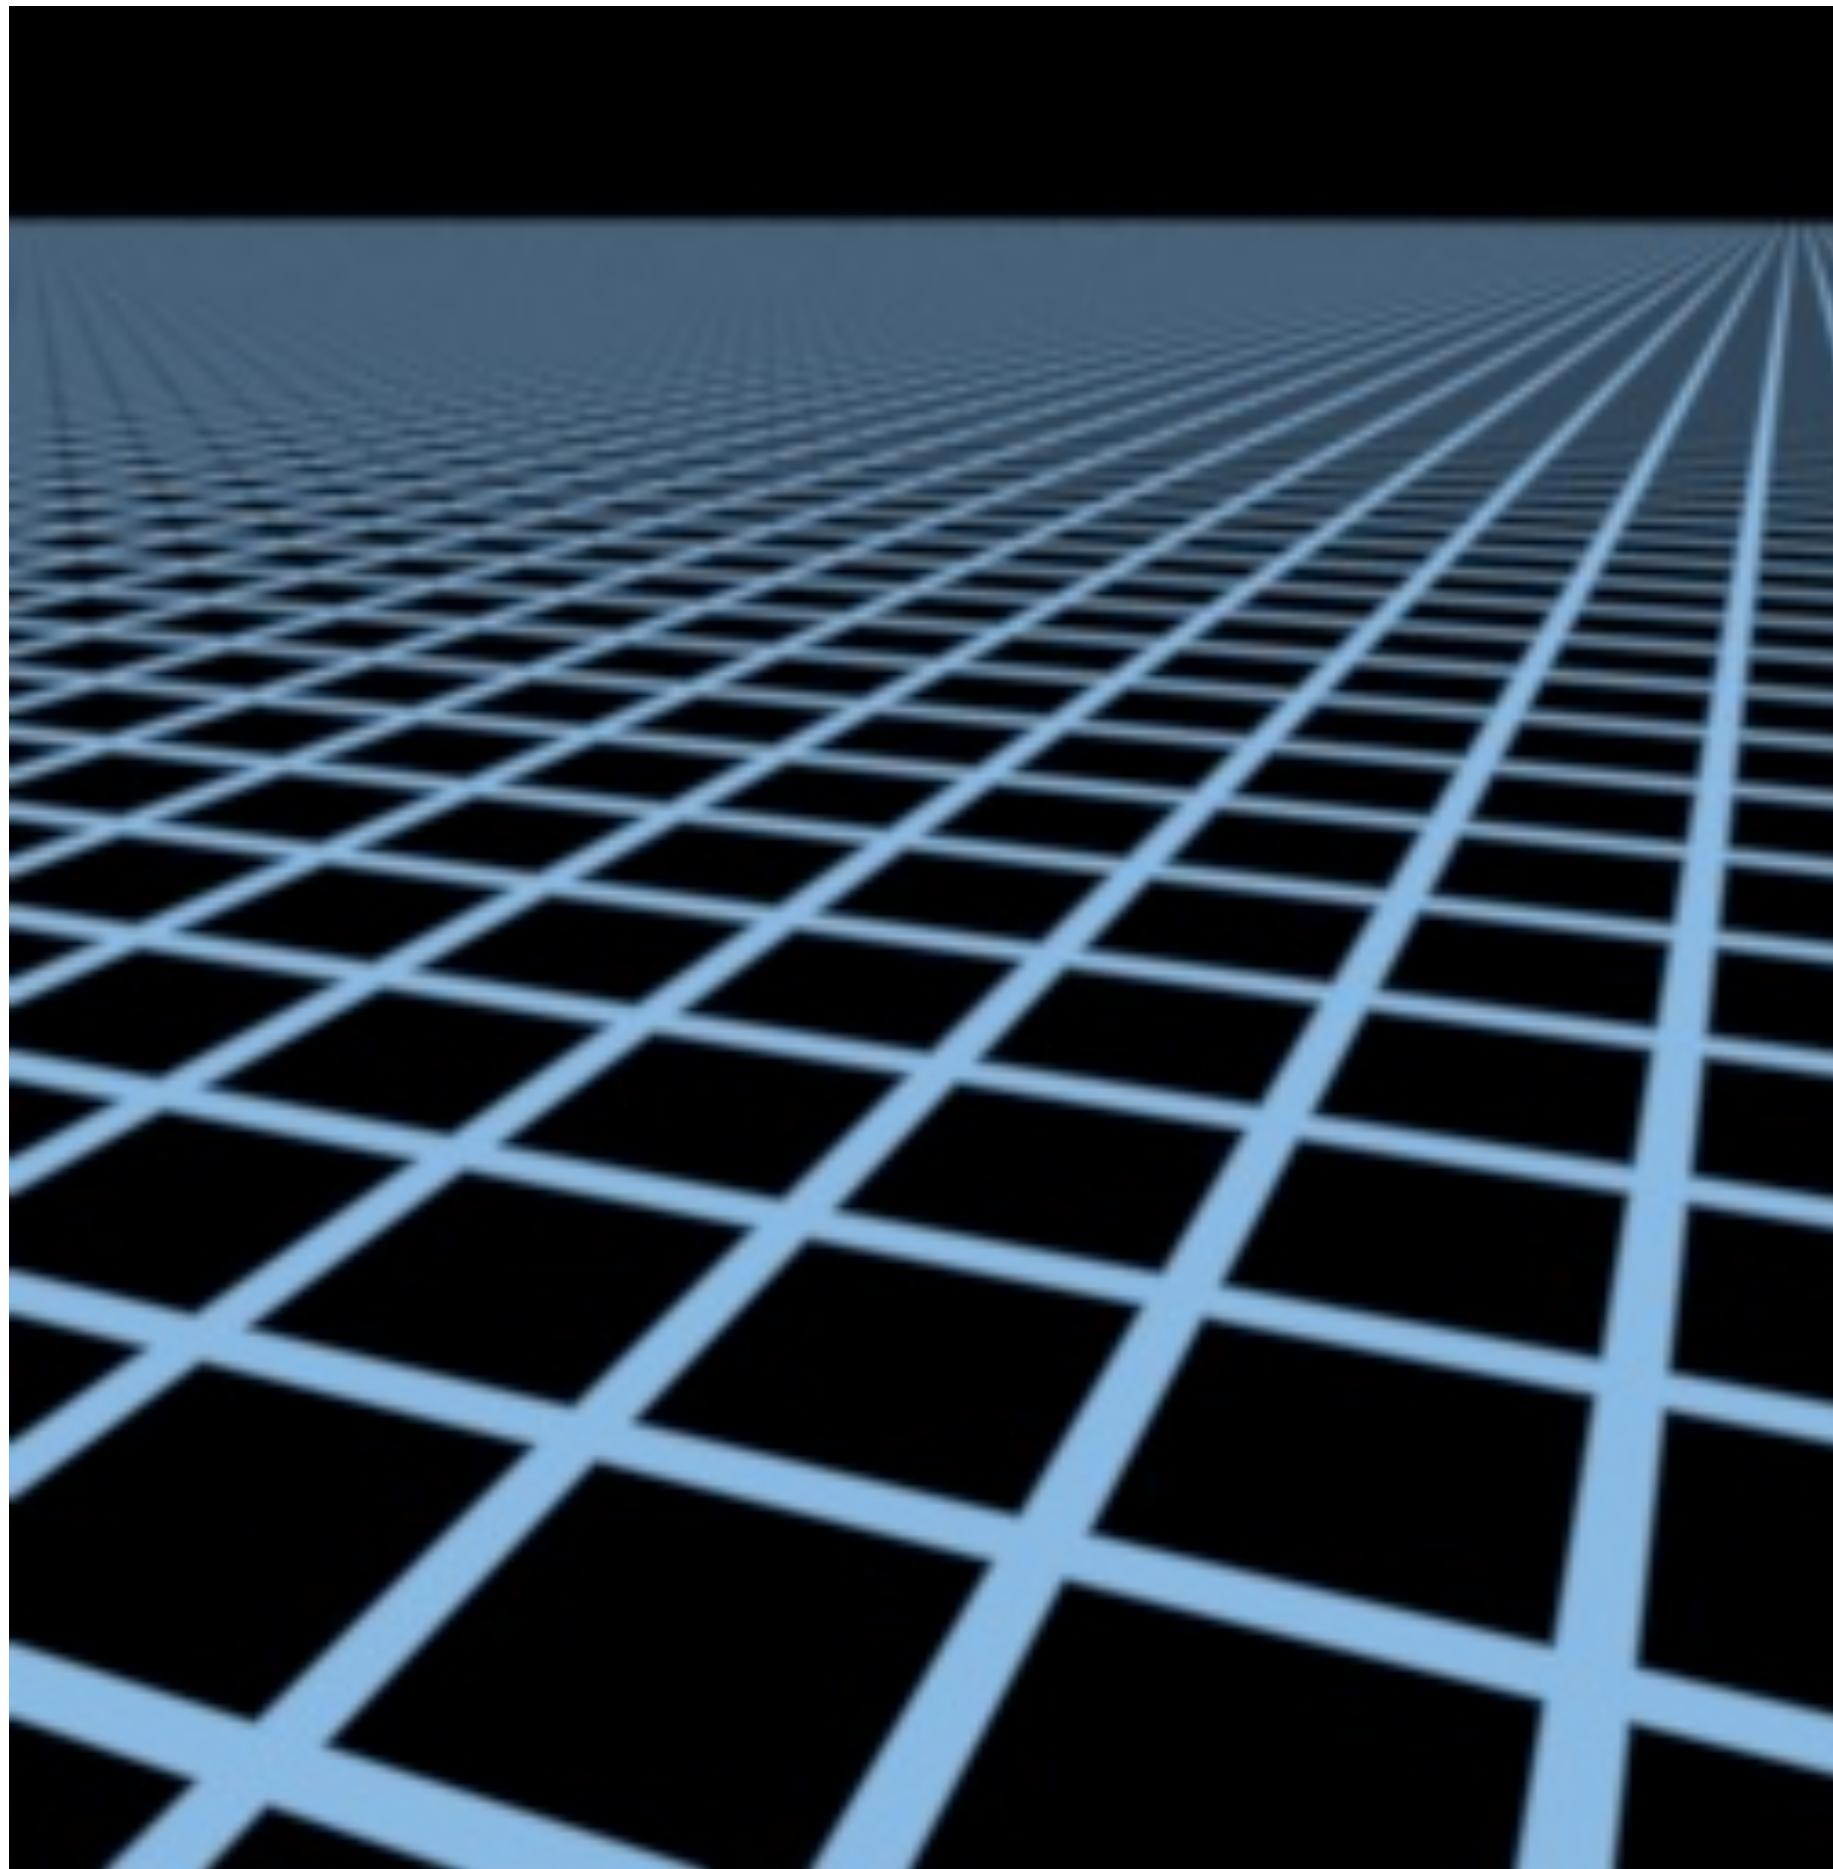
■ Bilinear resampling:

- 4 texel reads
- 3 lerps (3 mul + 6 add)

■ Trilinear resampling:

- 8 texel reads
- 7 lerps (7 mul + 14 add)

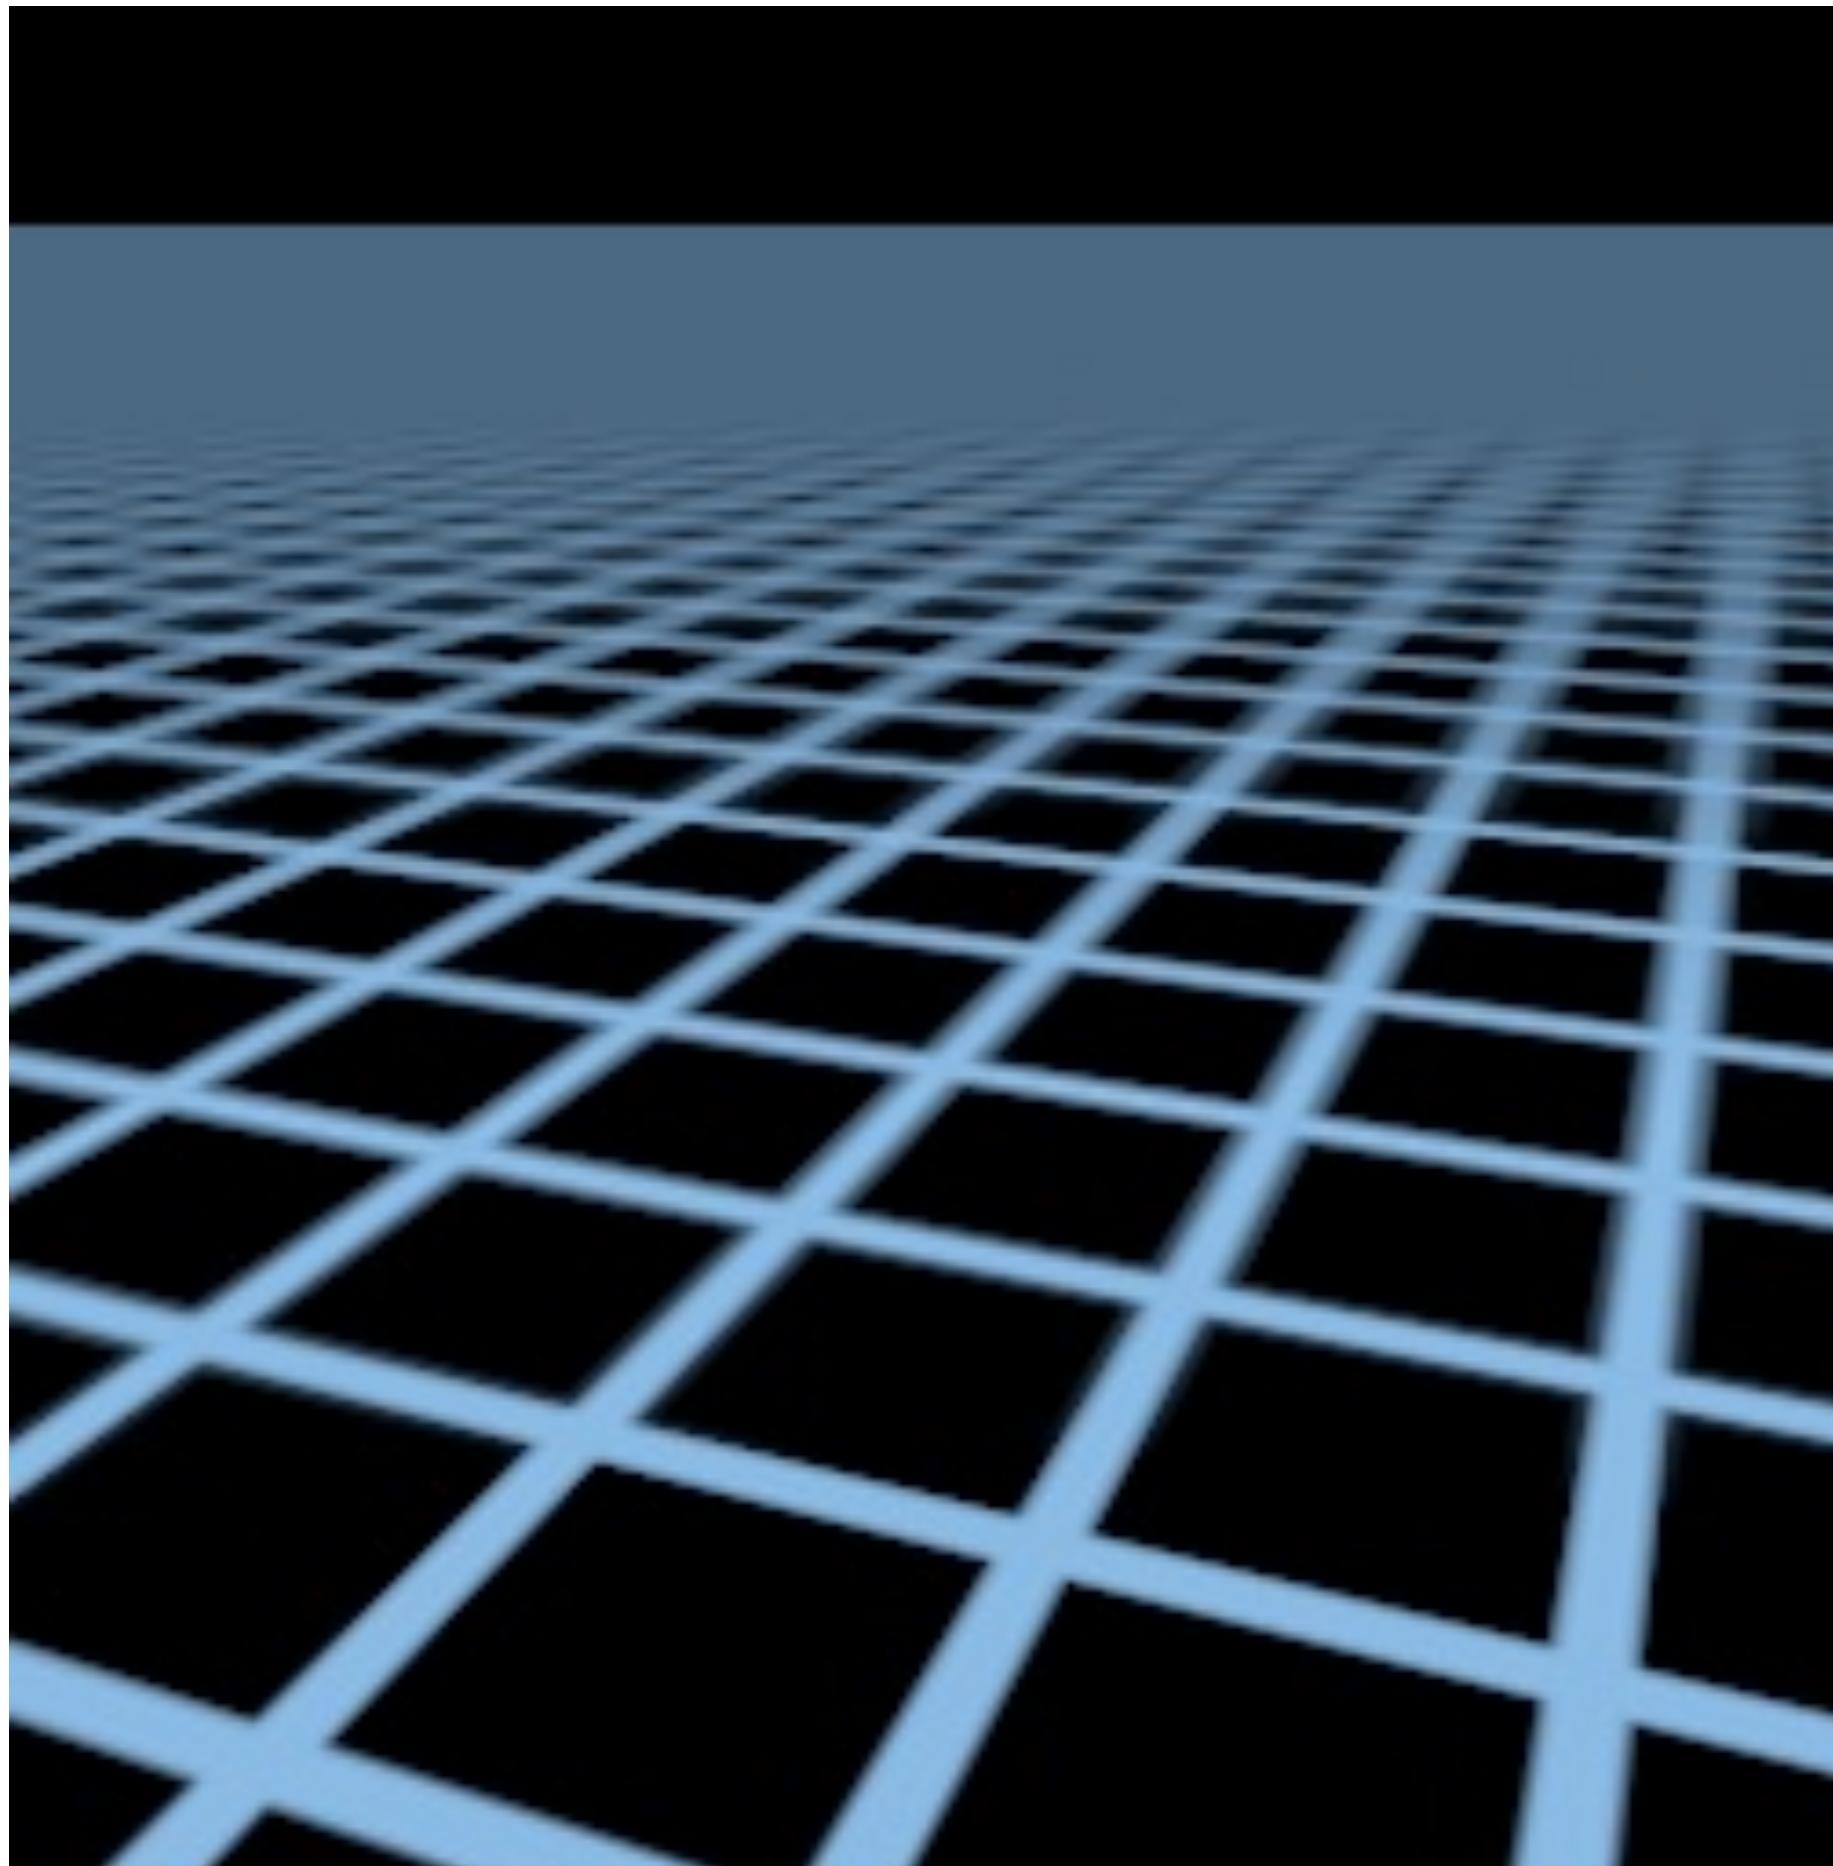
Example: mipmap limitations



**Supersampling: 512 texture samples per pixel
(desired answer)**

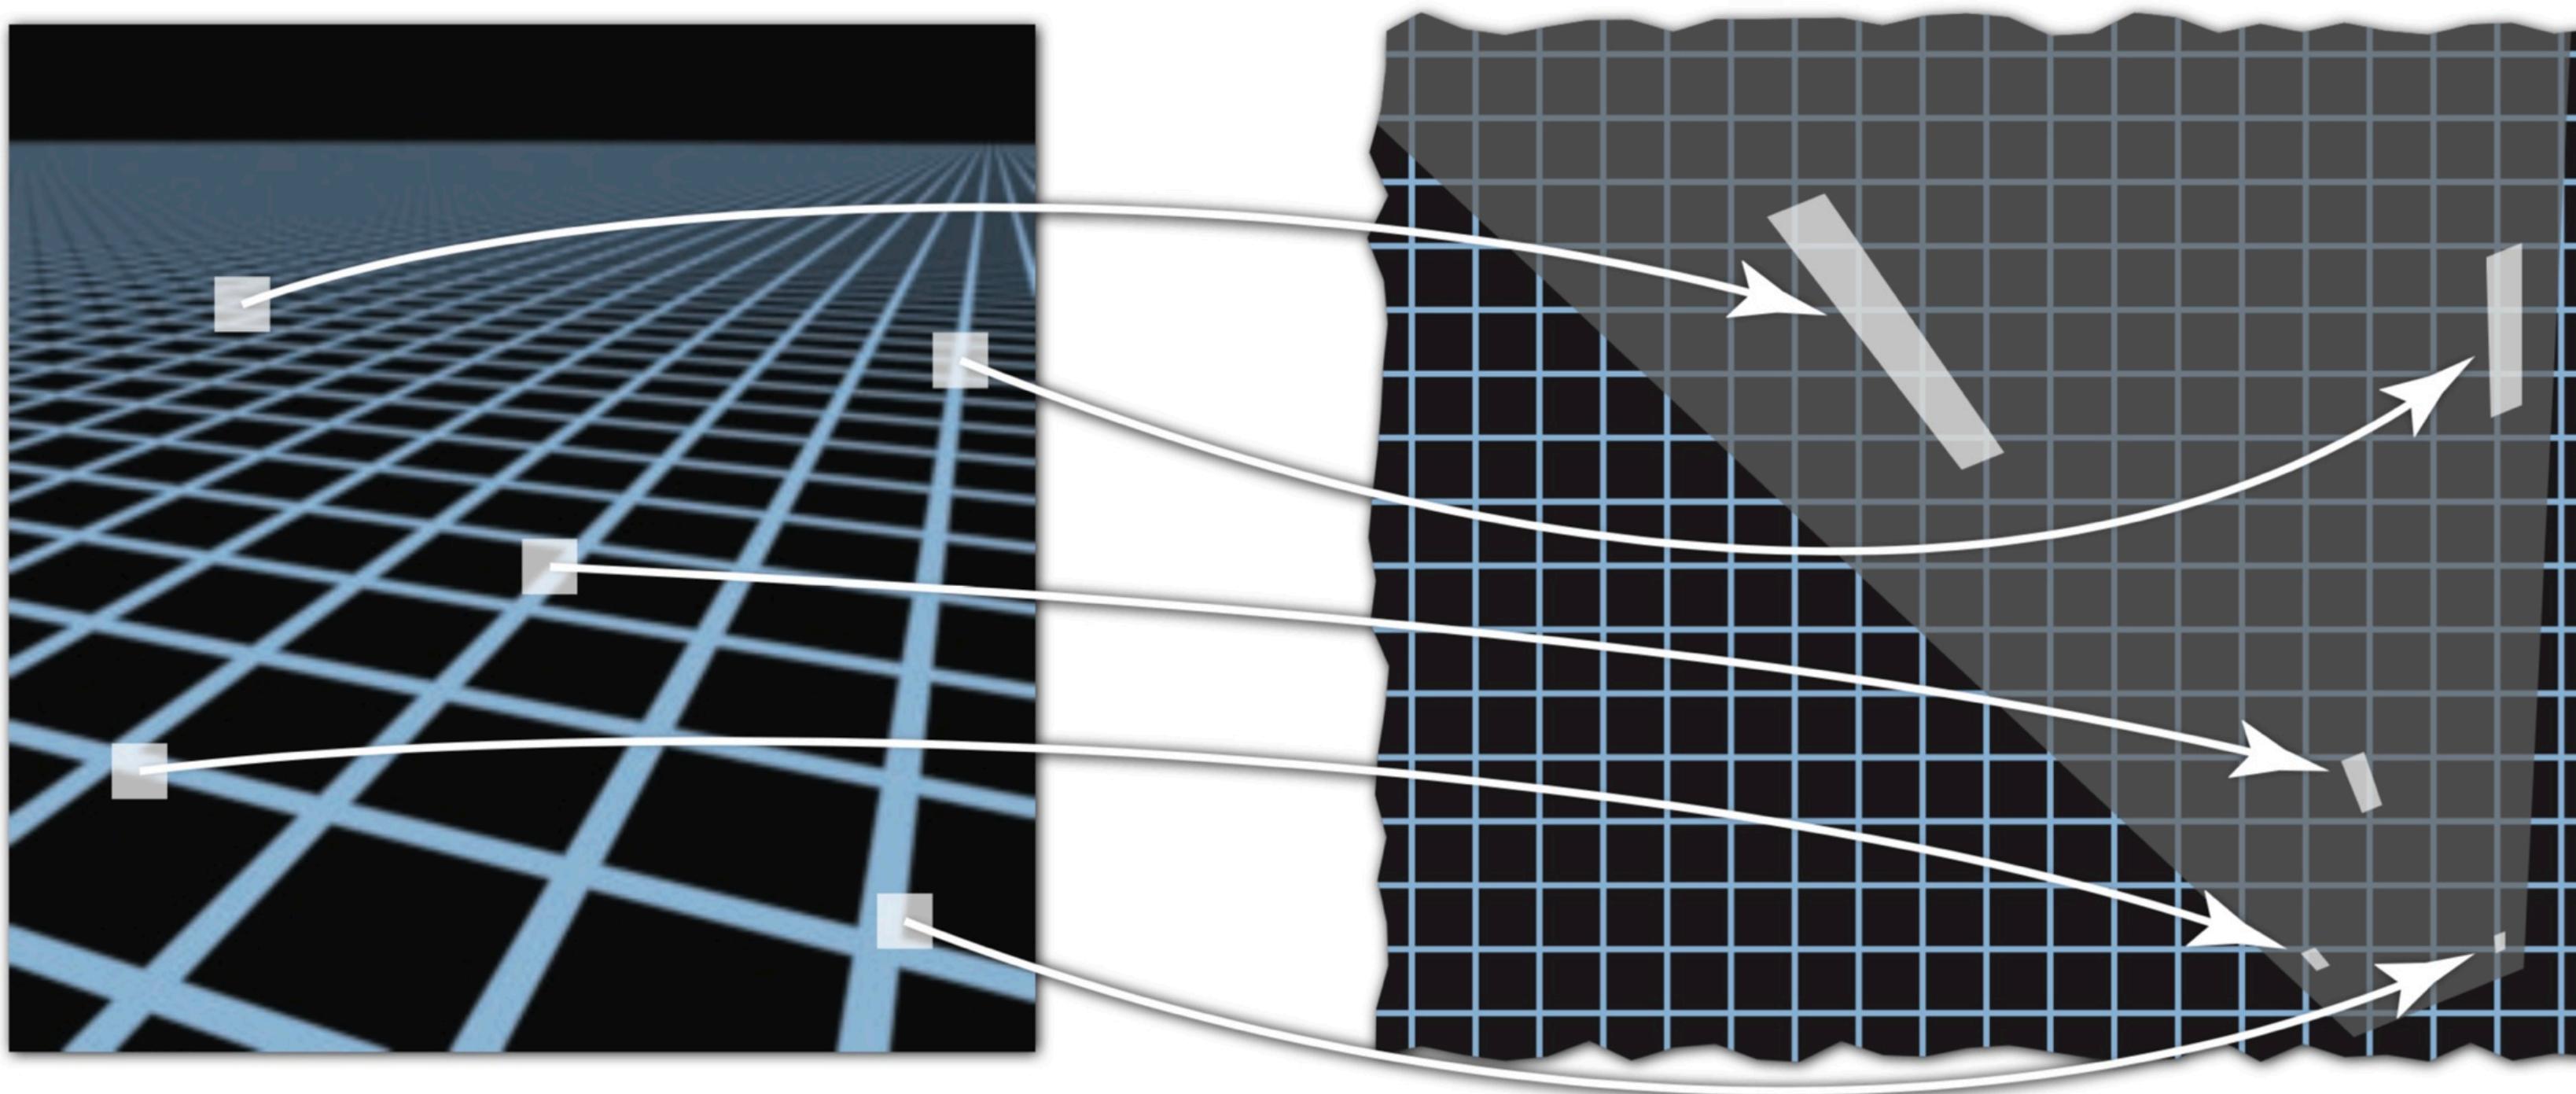
Example: mipmap limitations

Overblurs
Why?



Mipmap trilinear sampling

Screen pixel footprint in texture space



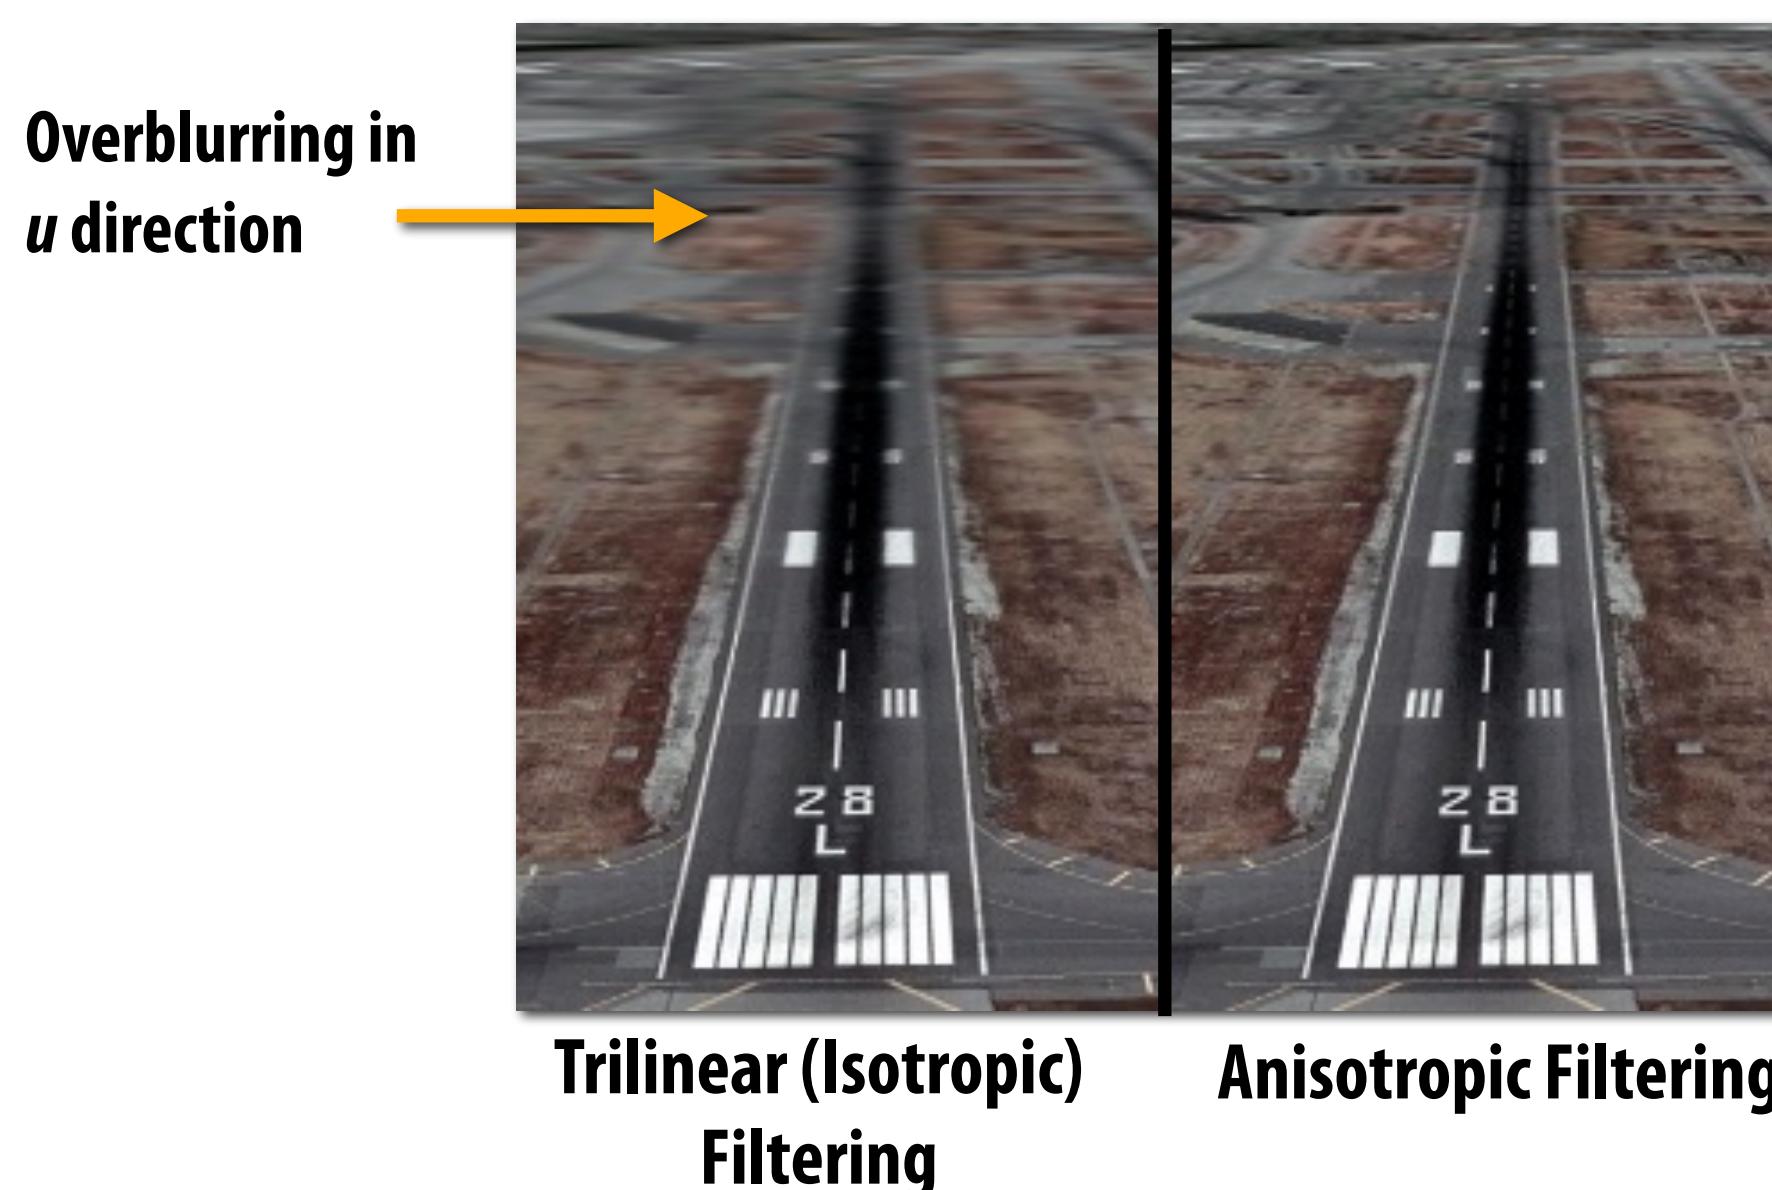
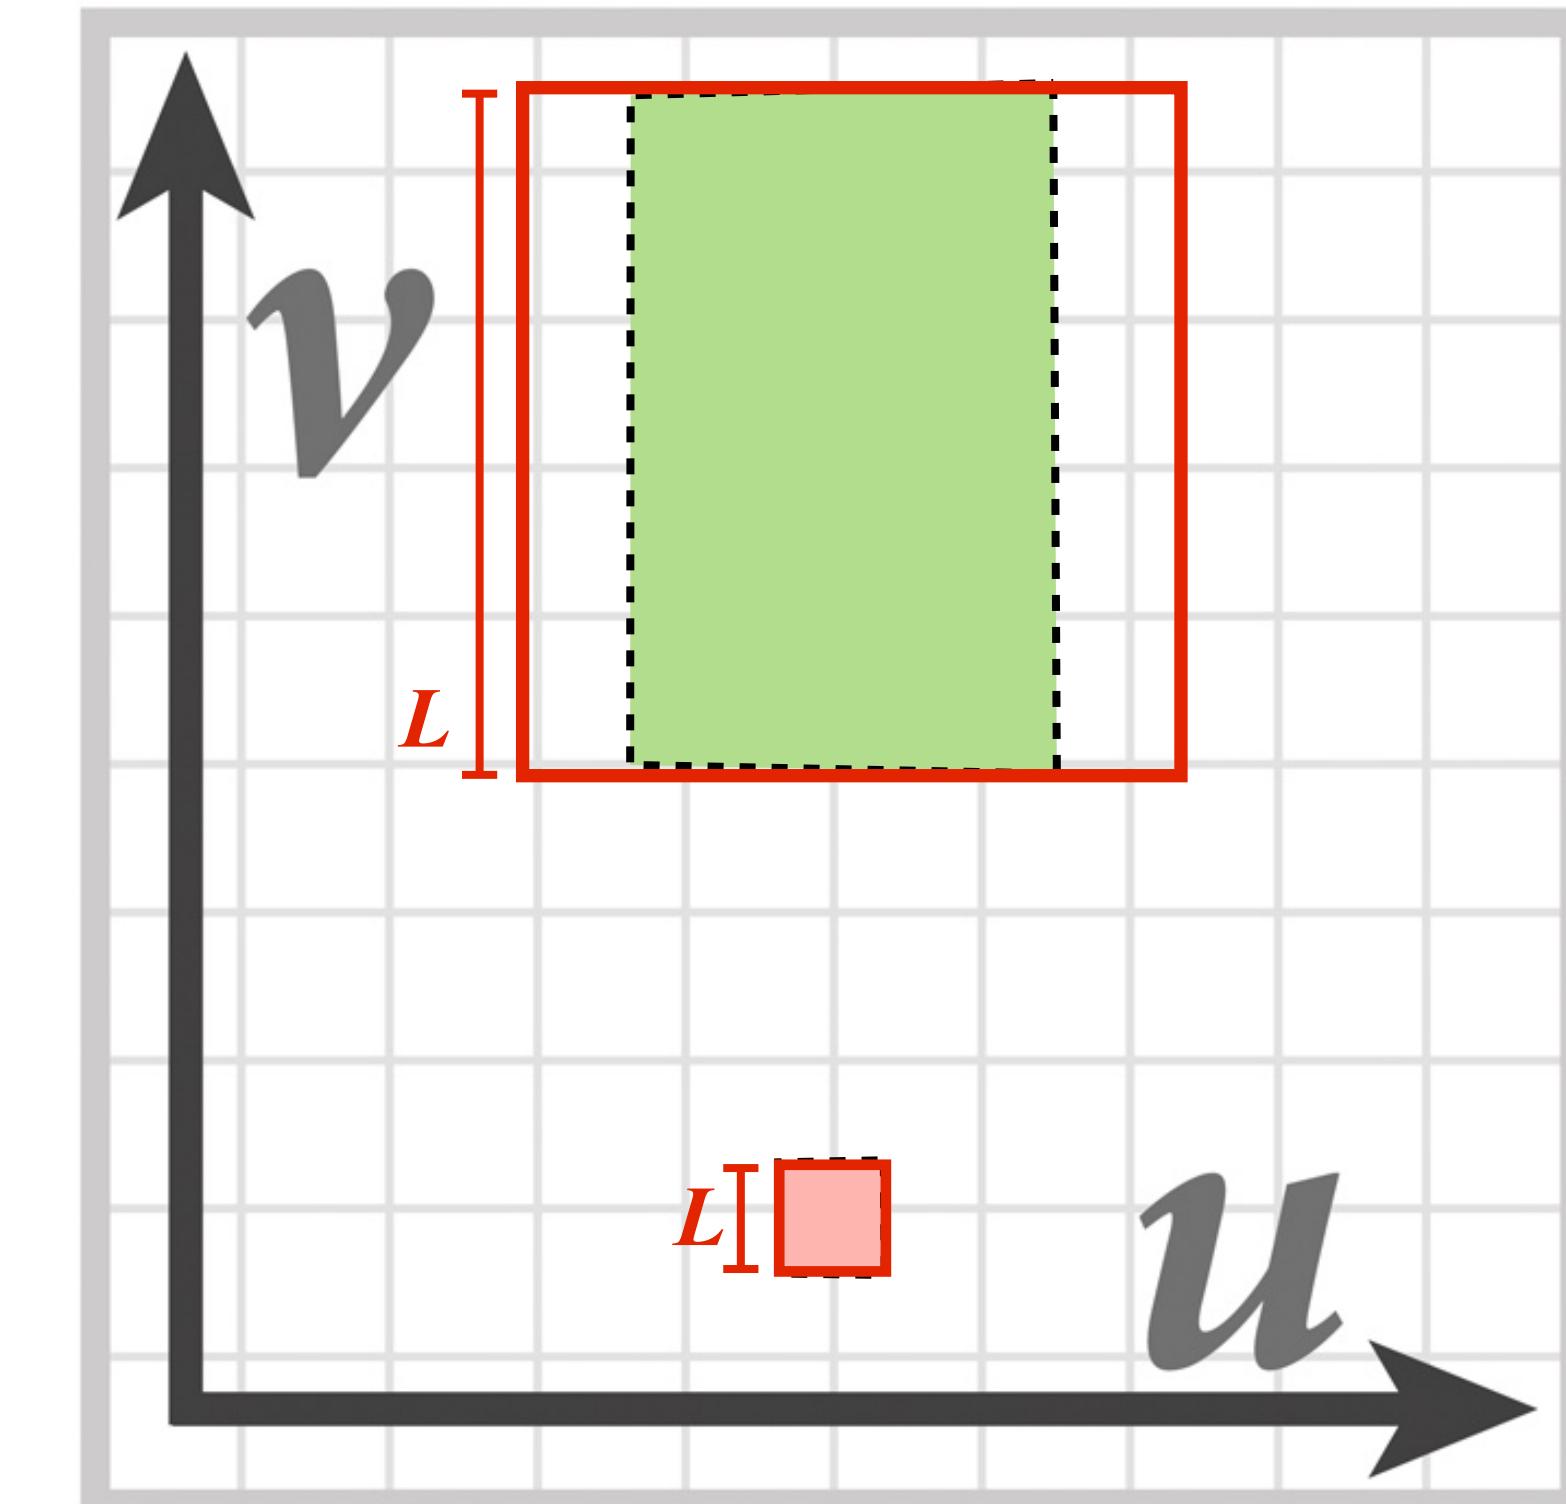
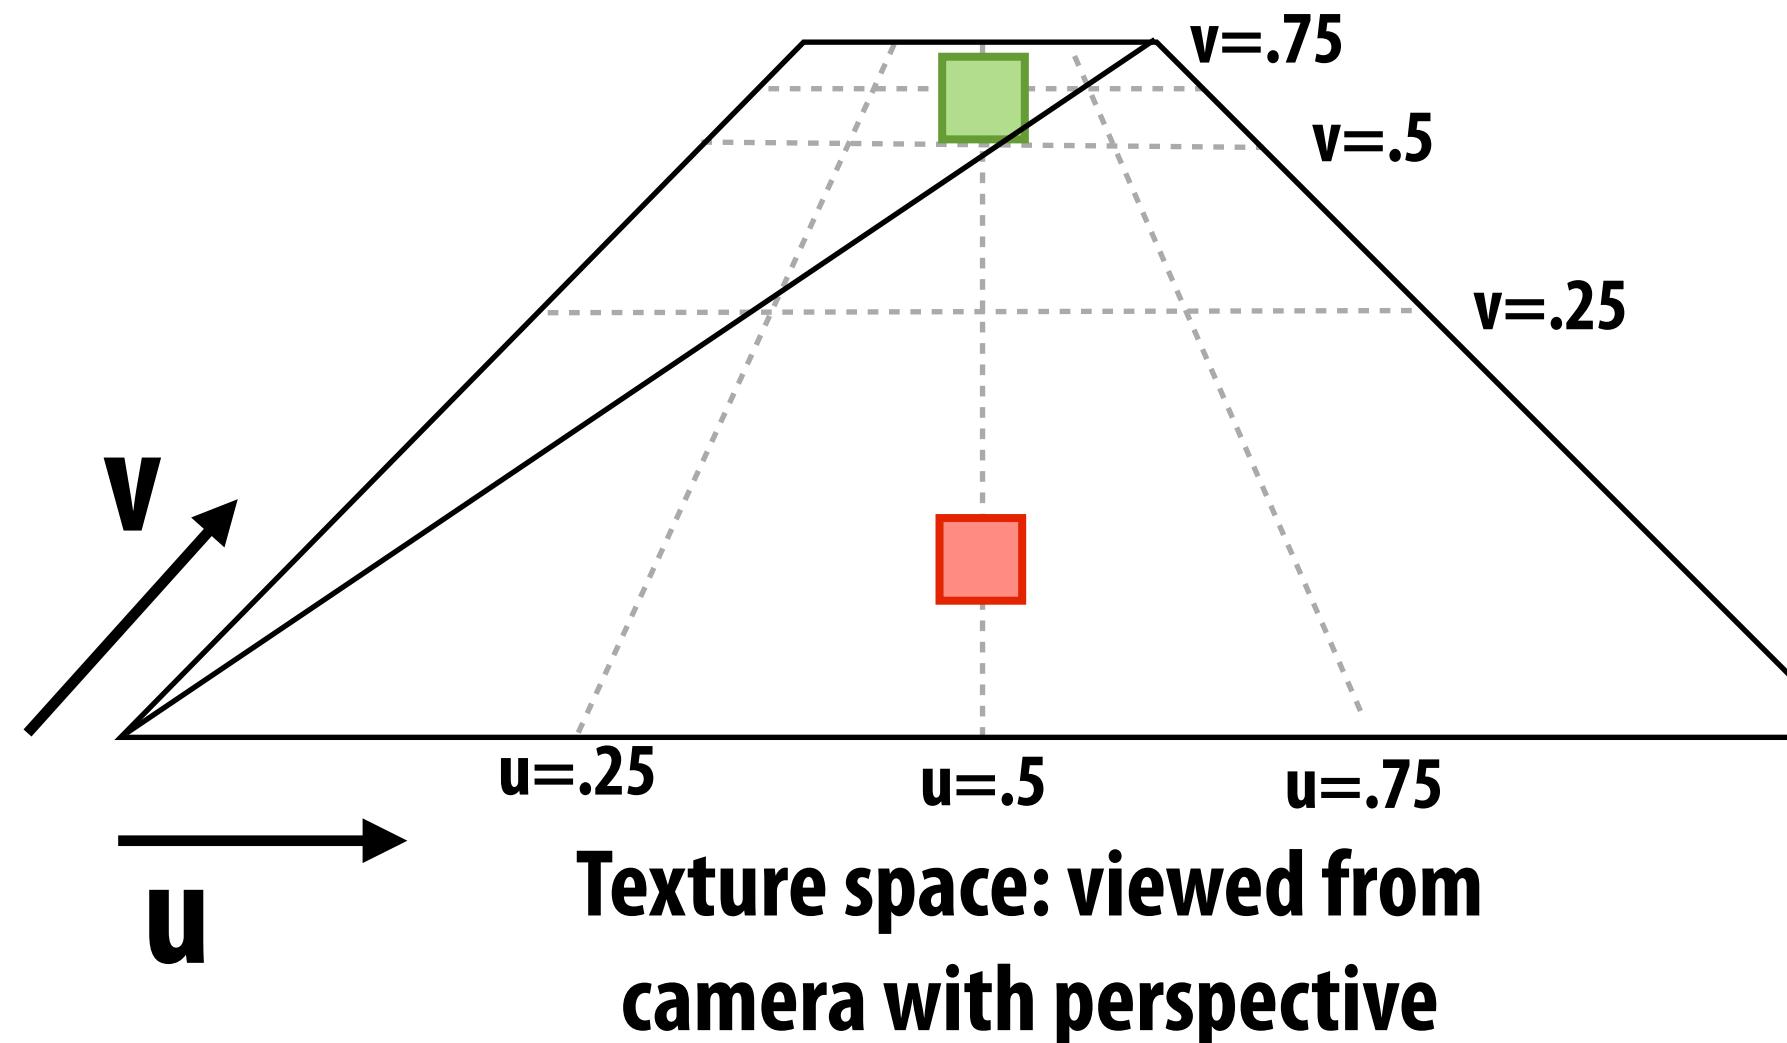
Screen space

Texture space

Texture sampling pattern not rectilinear or isotropic

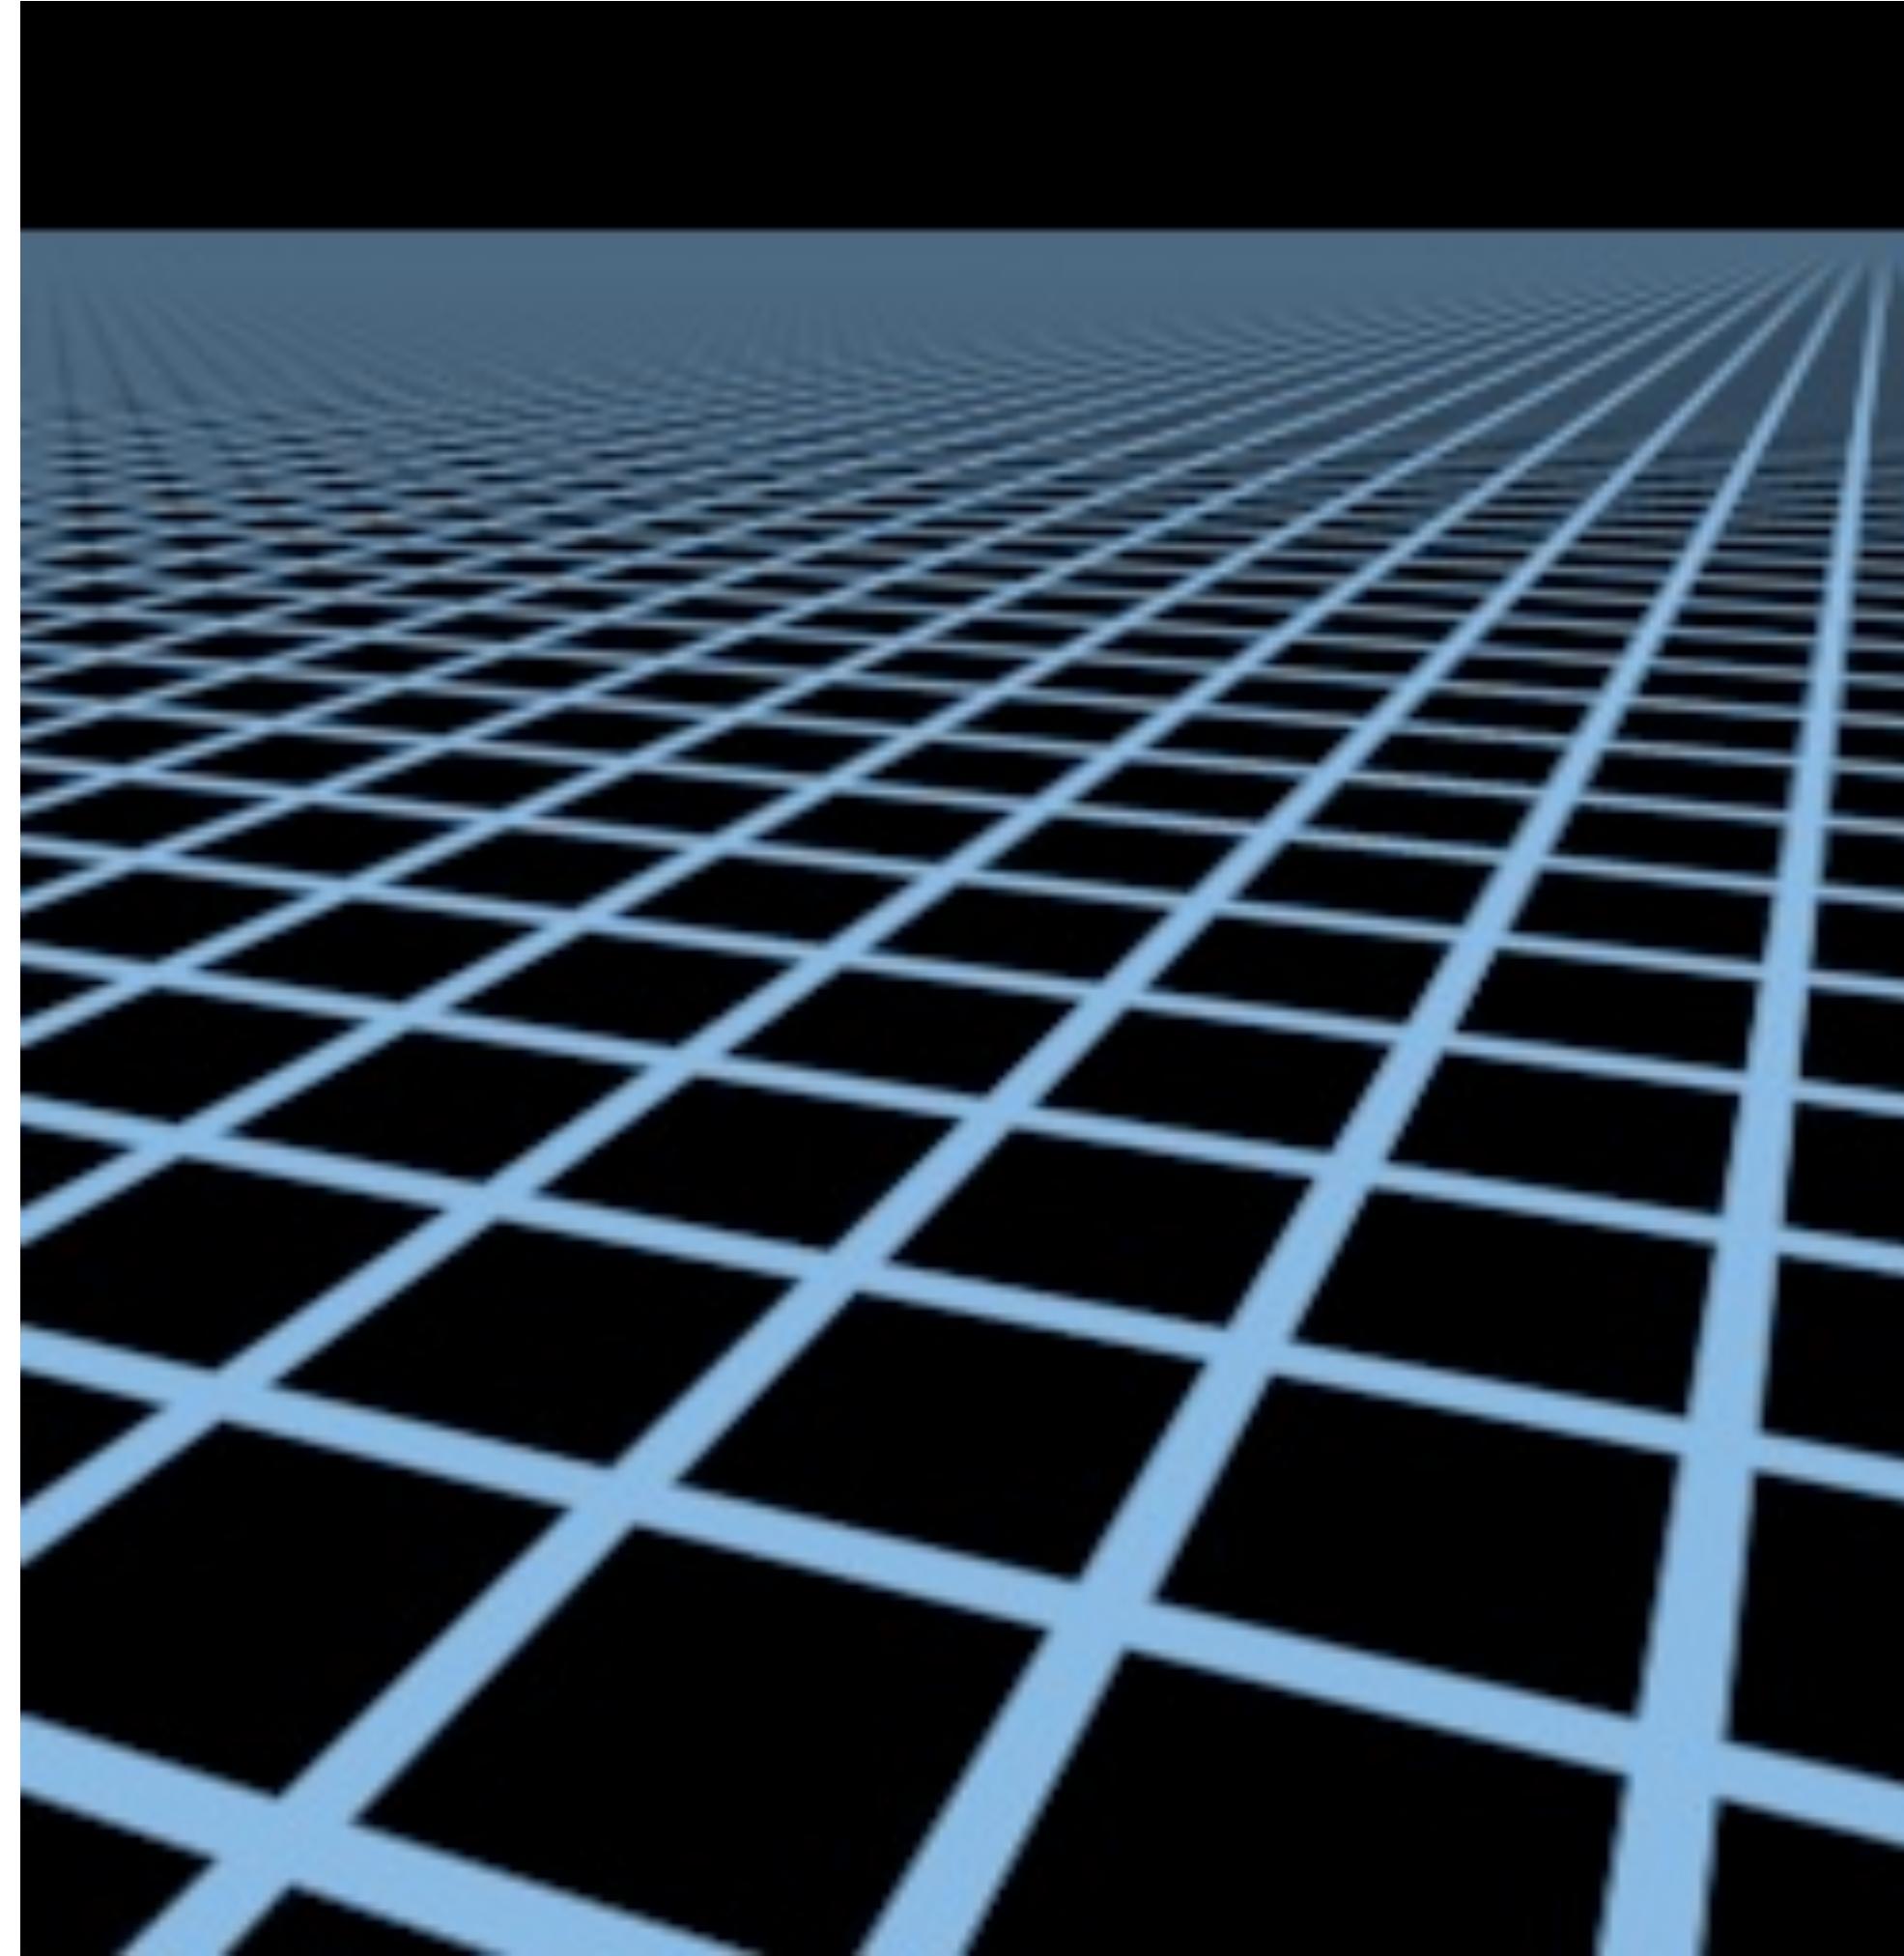
Pixel area may not map to isotropic region in texture space

Proper filtering requires anisotropic filter footprint



(Modern anisotropic texture filtering solutions combine multiple mip map samples to approximate integral of texture value over arbitrary texture space regions)

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 (in a ray caster, that's a ray-triangle test)
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.

Summary: texture mapping

- **Texturing: used to add visual detail to surfaces that is not captured in geometry**
- **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: precompute and store multiple multiple resampled versions of the texture image (each with different amounts of low-pass filtering to remove increasing amounts of high frequency detail)**
 - **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**

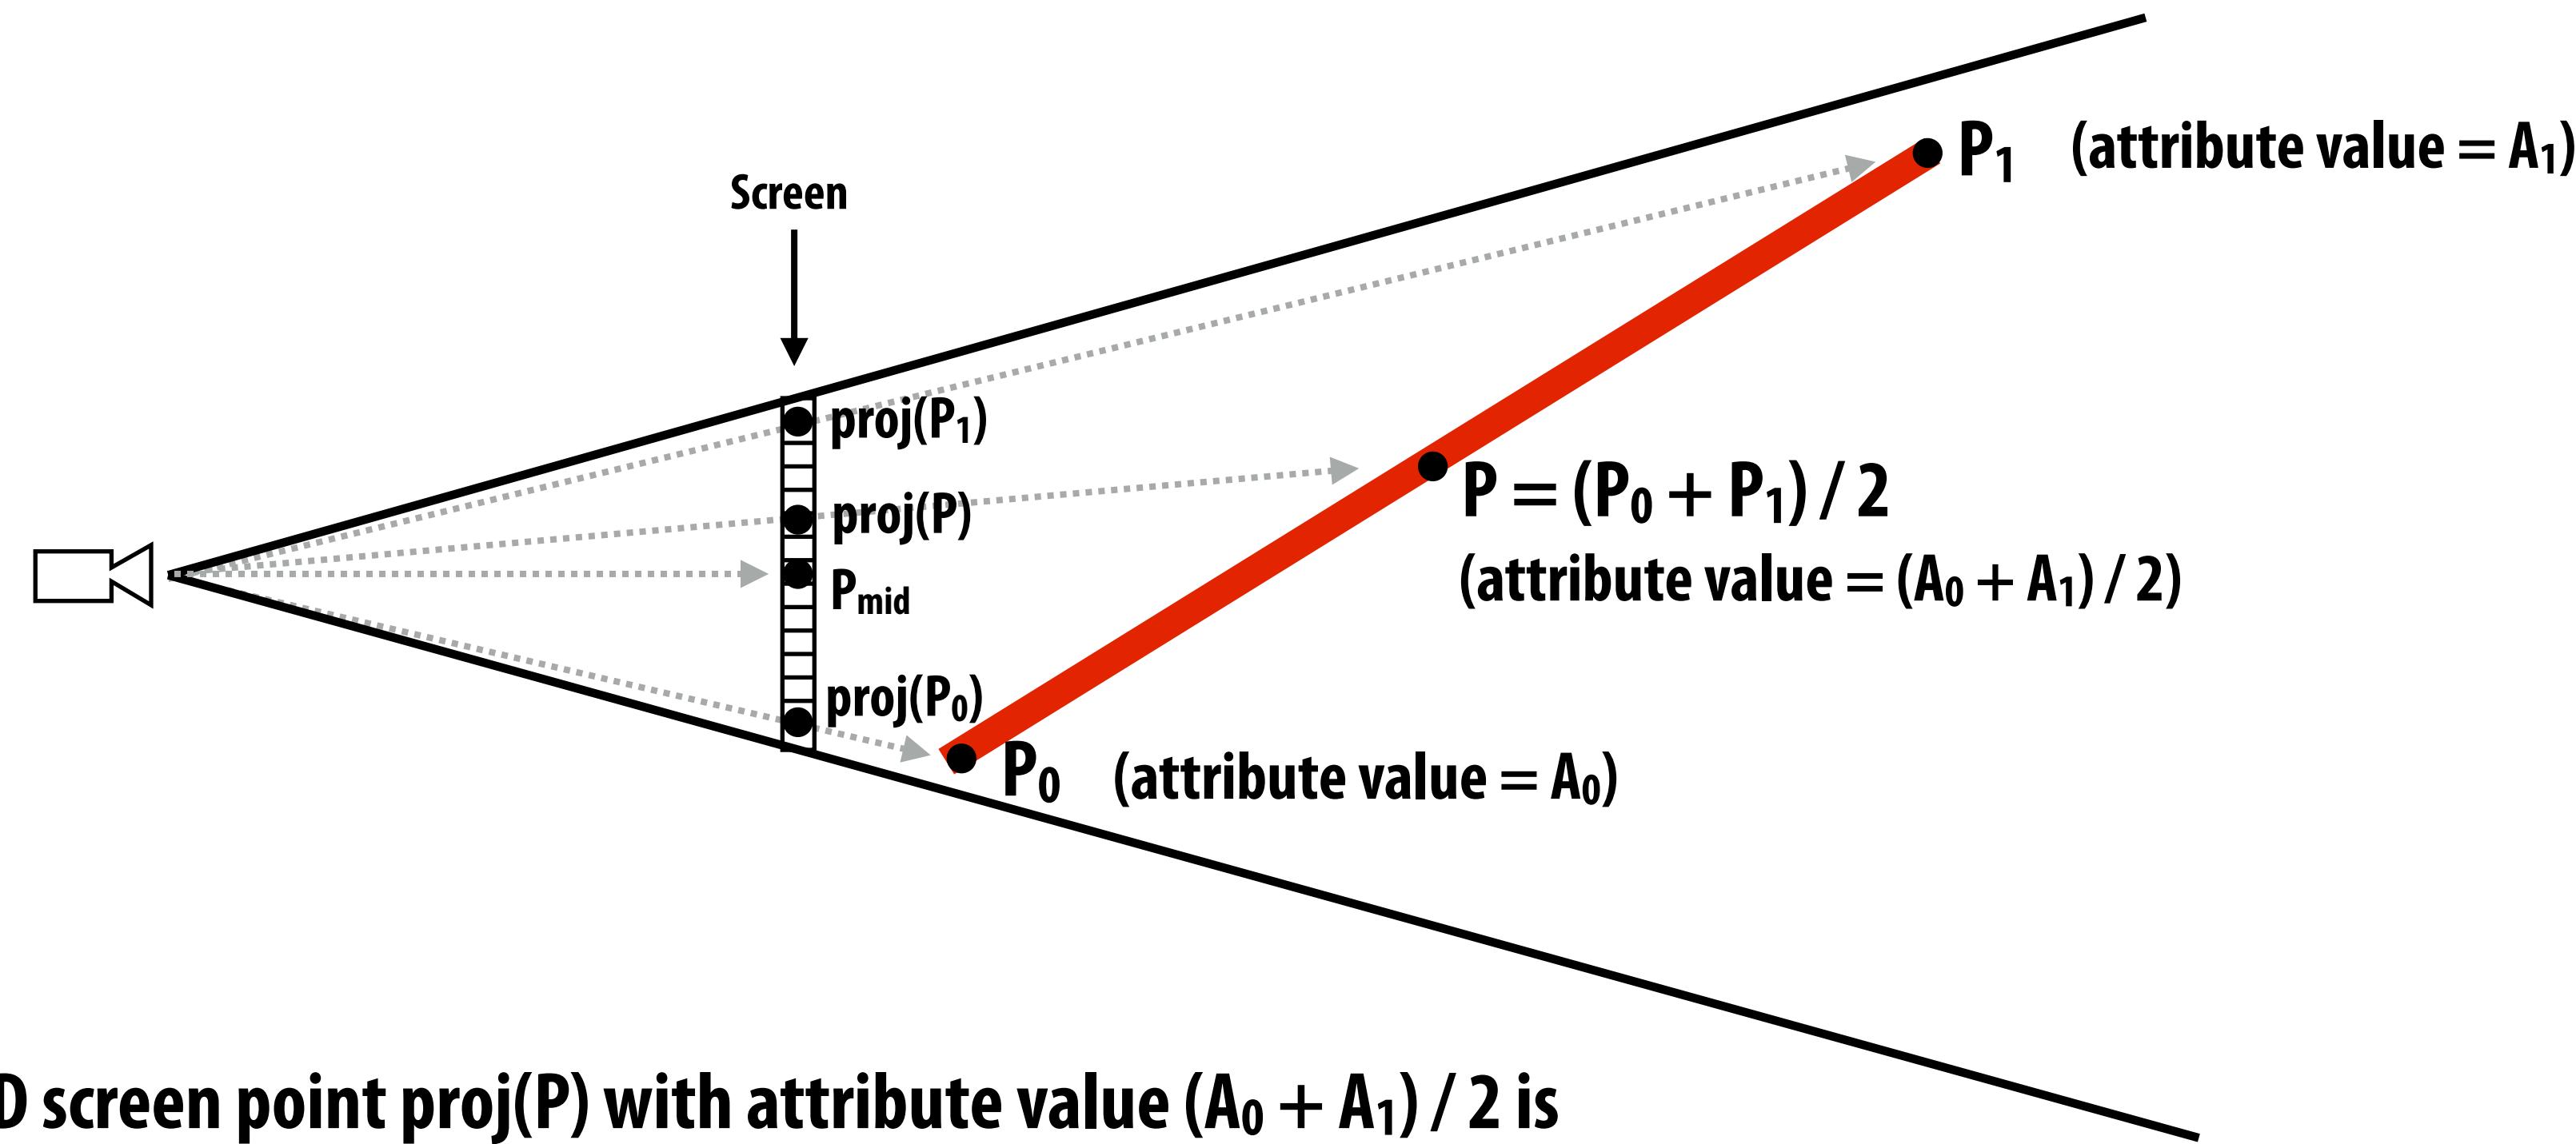
Bonus material (not covered in class)

Perspective correct interpolation in 2D

Not so fast... perspective incorrect interpolation

The value of an attribute at the 3D point P on a triangle is a linear combination of attribute values at vertices.

But due to perspective projection, barycentric interpolation of values on a triangle with vertices of different depths in 3D is not linear in 2D screen XY coordinates (vertex coordinates **after** projection)

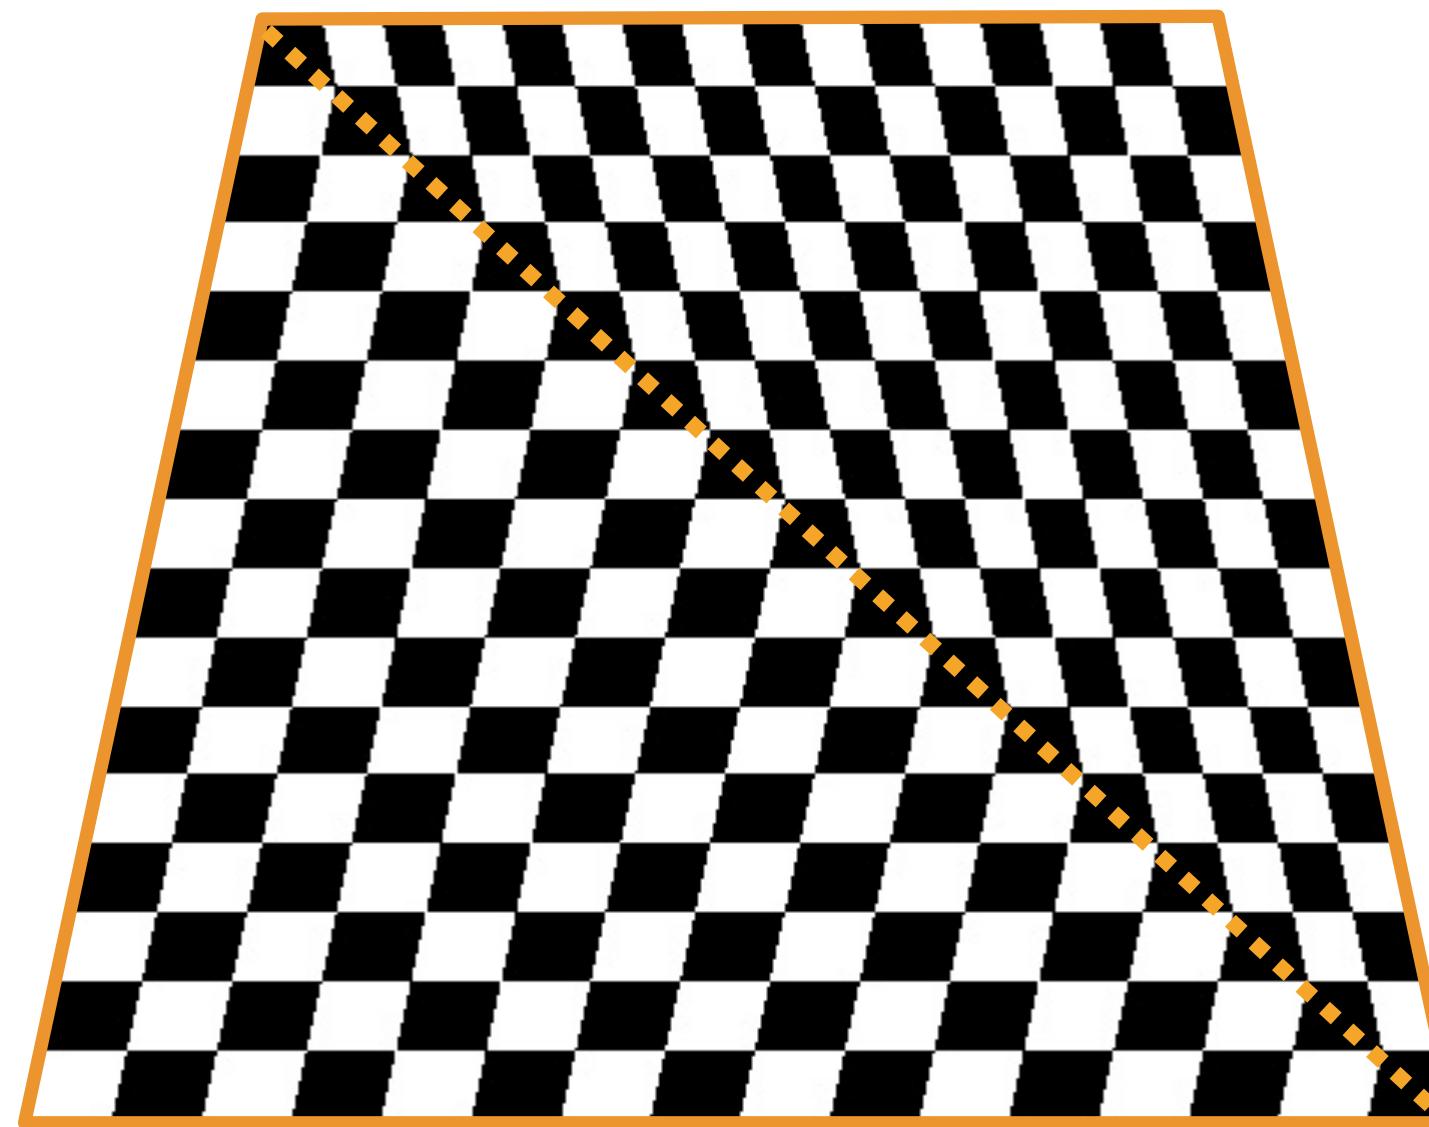


In this example, the 2D screen point $proj(P)$ with attribute value $(A_0 + A_1) / 2$ is not halfway between the 2D screen points $proj(P_0)$ and $proj(P_1)$.

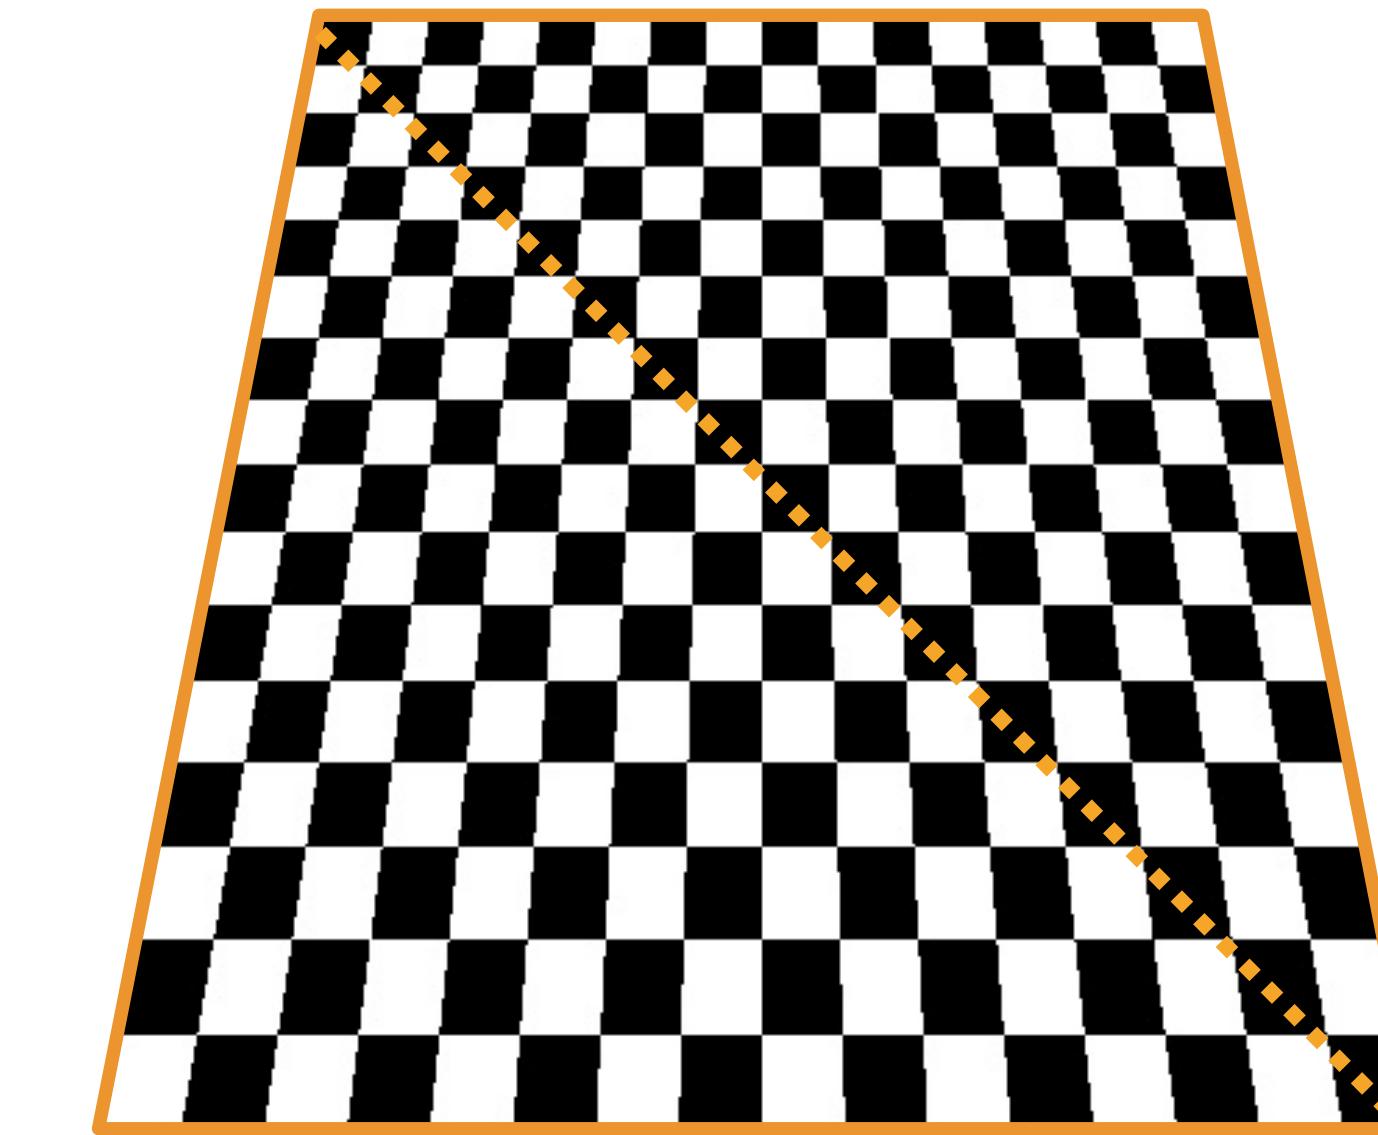
Similarly, the attribute's value at $P_{mid} = (proj(P_0) + proj(P_1)) / 2$ is not $(A_0 + A_1) / 2$.

Perspective correct interpolation

This is a plane (two triangles), tilted down and rendered under perspective.

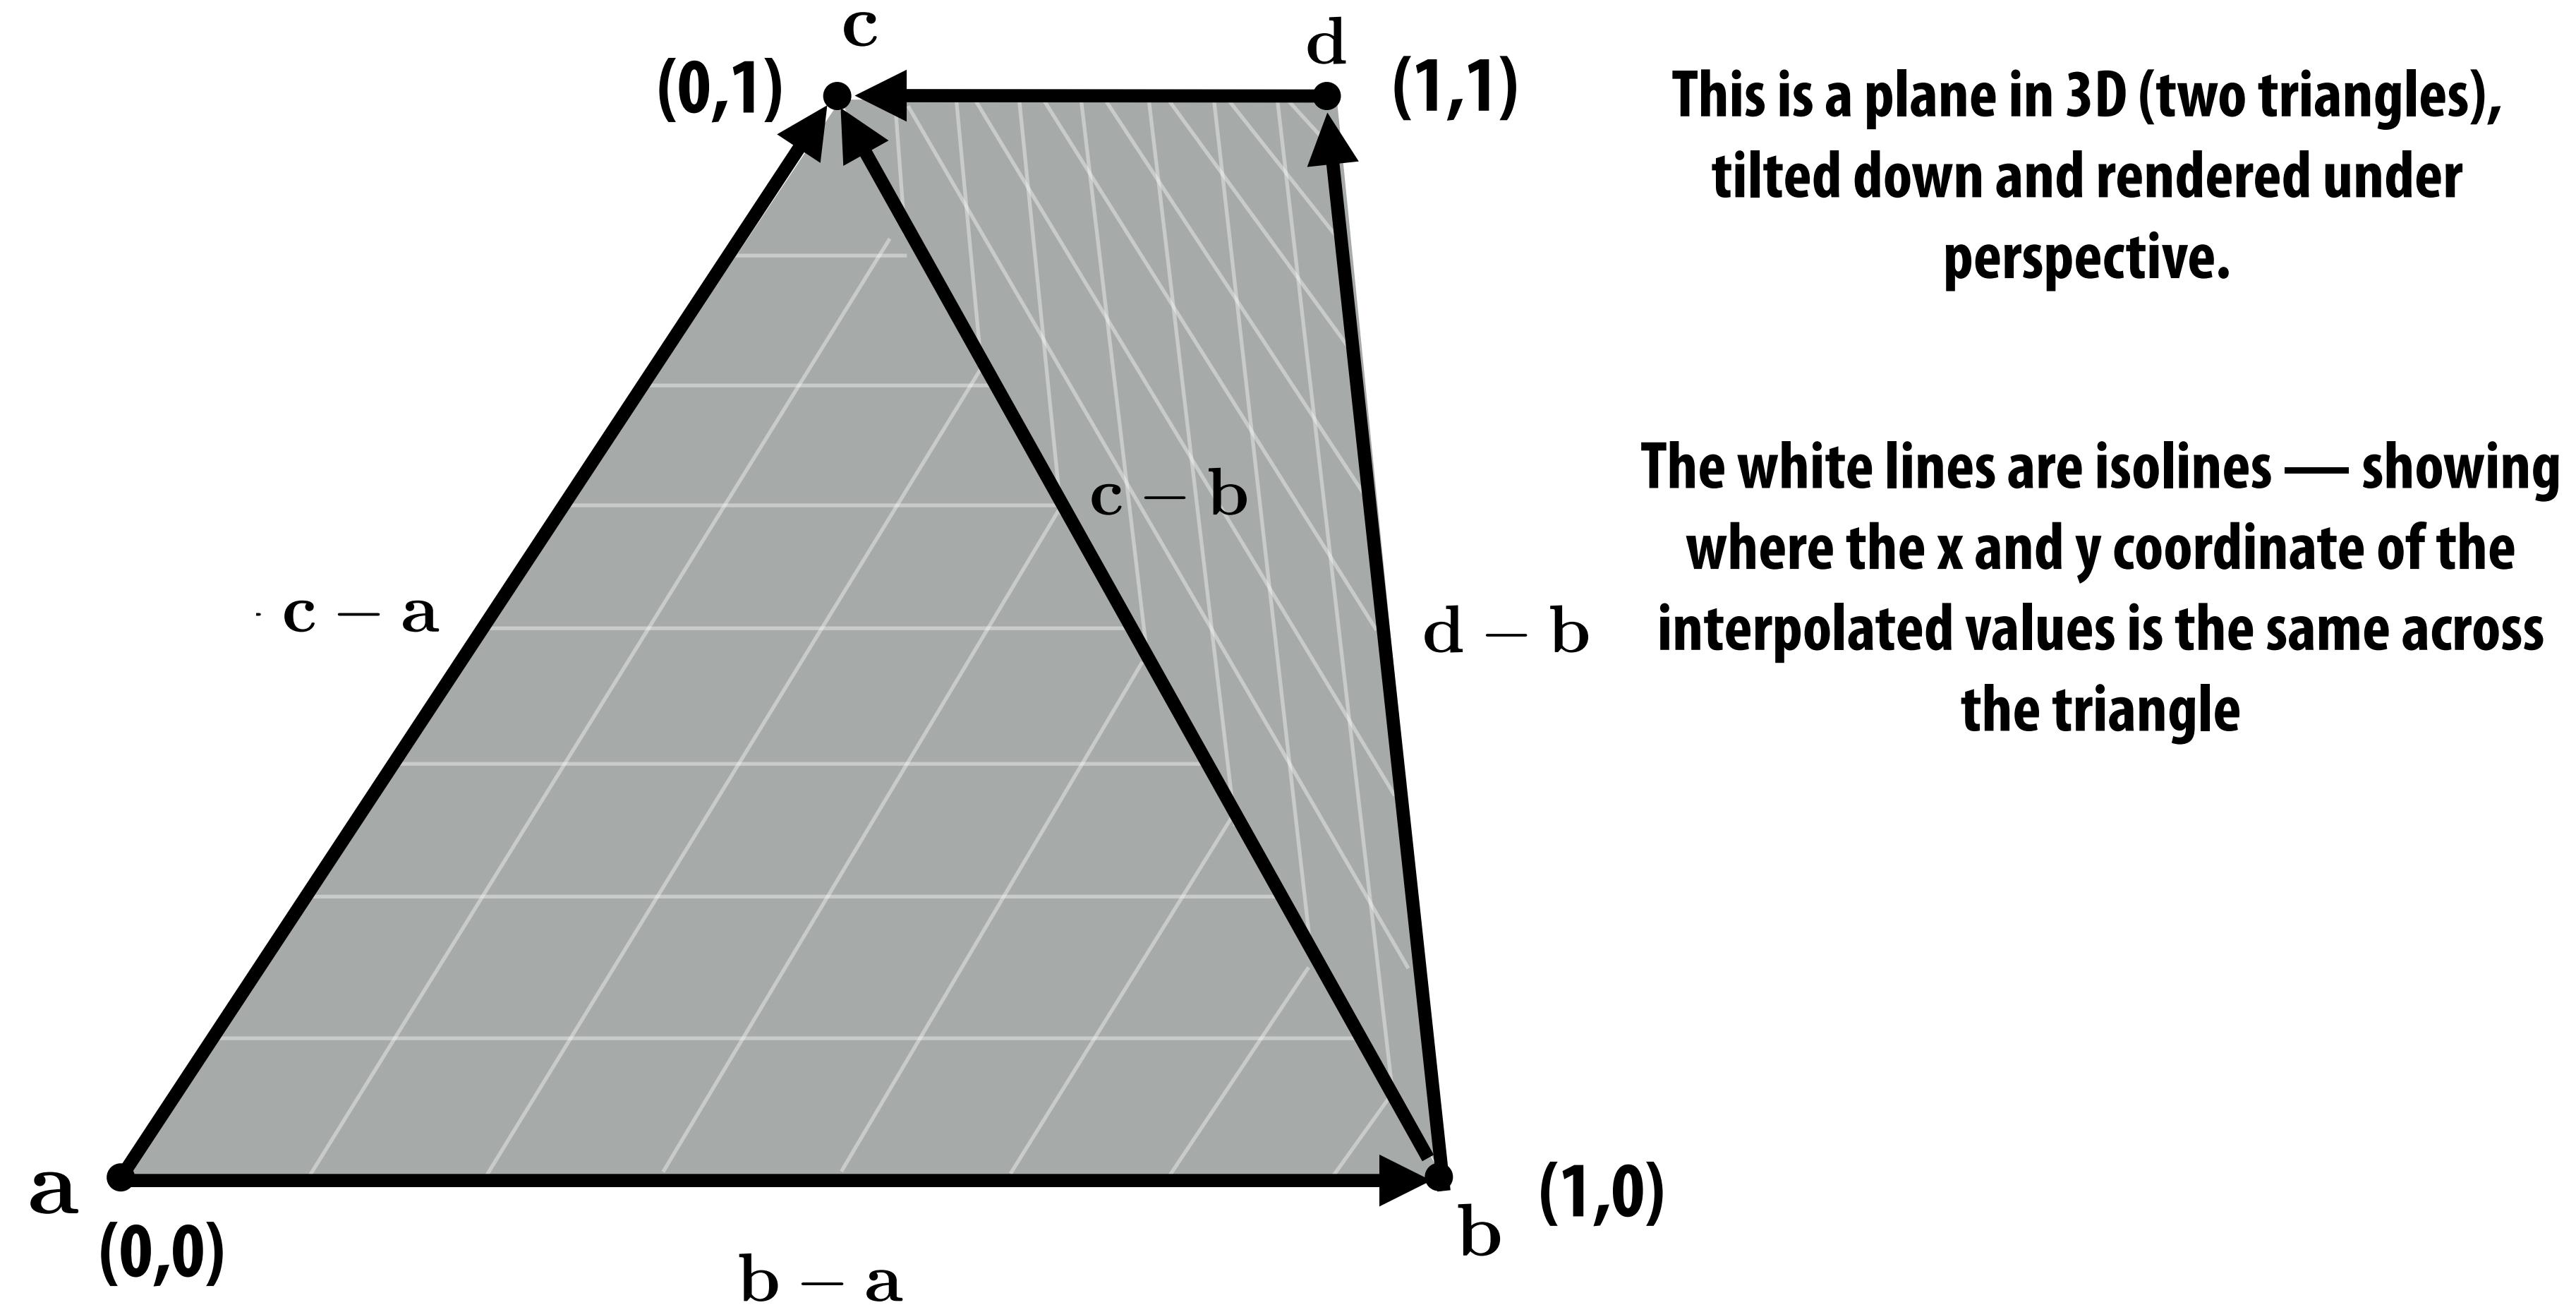


**2D screen-space
interpolation**



**3D world-space
interpolation**

Linear interpolation in screen space



Perspective correct interpolation on a projected triangle (in 2D)

■ Problem:

- Given some value f_i at each of a 3D triangle's vertices, that is linearly interpolated across the triangle in 3D
- And the 2D screen coordinates $P_i = (x_i, y_i)$ of each of a triangle's vertices after projection
- As well as the homogenous coordinate w_i for each vertex

Sample the value of $f(x, y)$ for the projected triangle at a given 2D screen space location (x, y)

Perspective-correct interpolation

Assume a triangle attribute varies linearly across the triangle (in 3D)

Attribute's value at 3D point on triangle $P = [x \ y \ z]^T$ is given by:

$$f(x, y, z) = ax + by + cz$$

Perspective project P , get 2D homogeneous representation:

$$\begin{bmatrix} x_{2D-H} \\ y_{2D-H} \\ w \end{bmatrix} = \begin{bmatrix} x \\ y \\ z \\ z \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix}$$

projection of P in 2D-H Drop z to move to 2D-H Simple perspective projection matrix * point P in 3D-H

perspective projection of P in 3D-H

* Note: using a more general perspective projection matrix only changes the coefficient in front of x_{2D} and y_{2D} . (property that f/w is affine still holds)

Then plug back in to equation for f at top of slide...

$$f(x_{2D-H}, y_{2D-H}) = ax_{2D-H} + by_{2D-H} + cw$$

$$\frac{f(x_{2D-H}, y_{2D-H})}{w} = \frac{a}{w}x_{2D-H} + \frac{b}{w}y_{2D-H} + c$$

$$\frac{f(x_{2D}, y_{2D})}{w} = \frac{a}{w}x_{2D} + \frac{b}{w}y_{2D} + c$$

So ... $\frac{f}{w}$ is affine function of 2D screen coordinates: $[x_{2D} \ y_{2D}]^T$

Direct evaluation of surface attributes from 2D-H vertices

For any surface attribute (with value defined at triangle vertices as: f_a, f_b, f_c)

w coordinate of vertex a after
perspective projection transform

$$\frac{f_a}{\mathbf{a}_w} = A\mathbf{a}_x + B\mathbf{a}_y + C$$

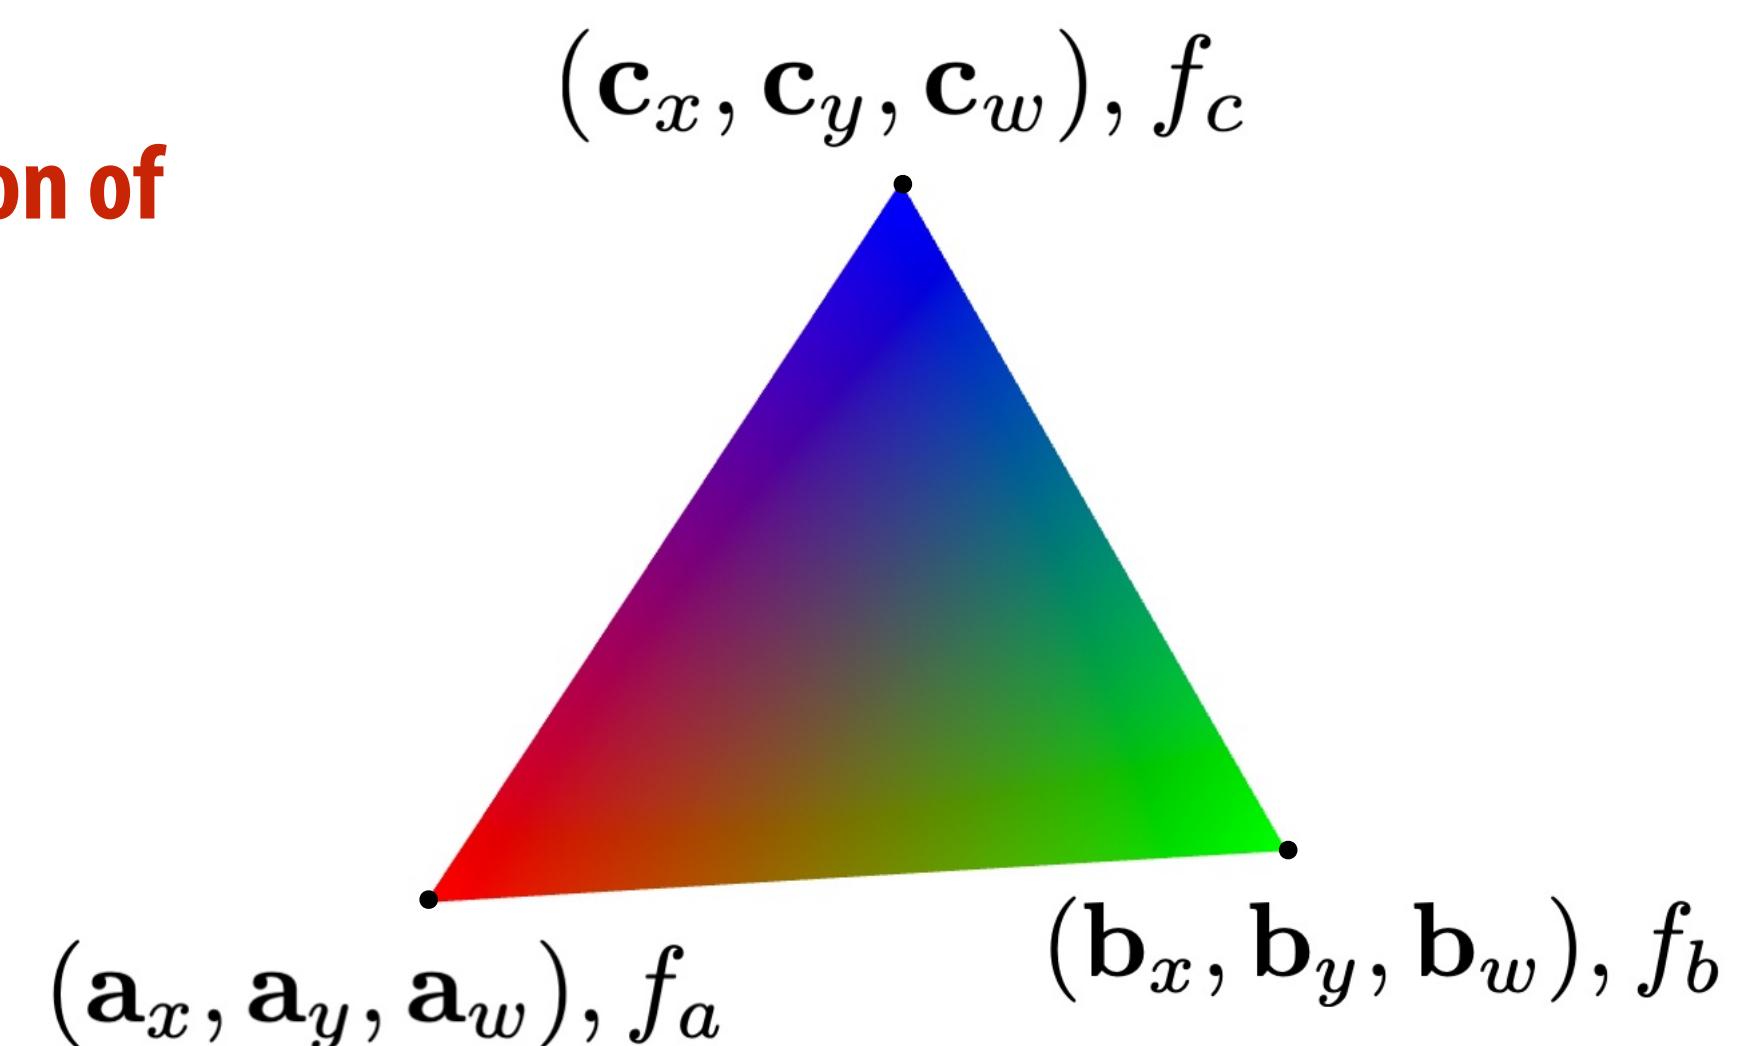
value of attribute at vertex a

projected 2D position of
vertex a

$$\frac{f_b}{\mathbf{b}_w} = A\mathbf{b}_x + B\mathbf{b}_y + C$$

$$\frac{f_c}{\mathbf{c}_w} = A\mathbf{c}_x + B\mathbf{c}_y + C$$

3 equations, solve for 3 unknowns (A, B, C)



This is done as a per triangle “setup” computation prior to sampling, just like you computed edge equations for evaluating coverage.

Efficient perspective-correct interpolation

Setup:

Given f_a, f_b, f_c and $w_a, w_b, w_c \dots$ compute A, B, C for $f/w(x,y) = Ax + By + C$

Also compute equation for $1/w(x,y)$

To evaluate surface attribute $f(x,y)$ at every covered sample (x,y) :

Evaluate $1/w(x,y)$

(from precomputed equation for value $1/w$)

Reciprocate $1/w(x,y)$ to get $w(x,y)$

For each triangle attribute:

Evaluate $f/w(x,y)$

(from precomputed equation for value f/w)

Multiply $f/w(x,y)$ by $w(x,y)$ to get $f(x,y)$

Works for any surface attribute f that varies linearly across triangle:

e.g., color, depth, texture coordinates

Acknowledgements

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