Filtering After Shading With Stochastic Texture Filtering

Matt Pharr, NVIDIA

Bartlomiej Wronski, NVIDIA

Marco Salvi, NVIDIA

Marcos Fajardo, Shiokara–Engawa Research

Intro – textures and texture filtering

(The way we traditionally teach it)

Textures – essential for high quality rendering

"Physically Based Rendering: From Theory To Implementation", 2004-2021 M. Pharr, W. Jakob, and G. Humphreys

Texture mapping – why we need filtering?

- Texture – a 1D/2D/3D/4D grid of discrete values
- Values defined only at **texel centers**

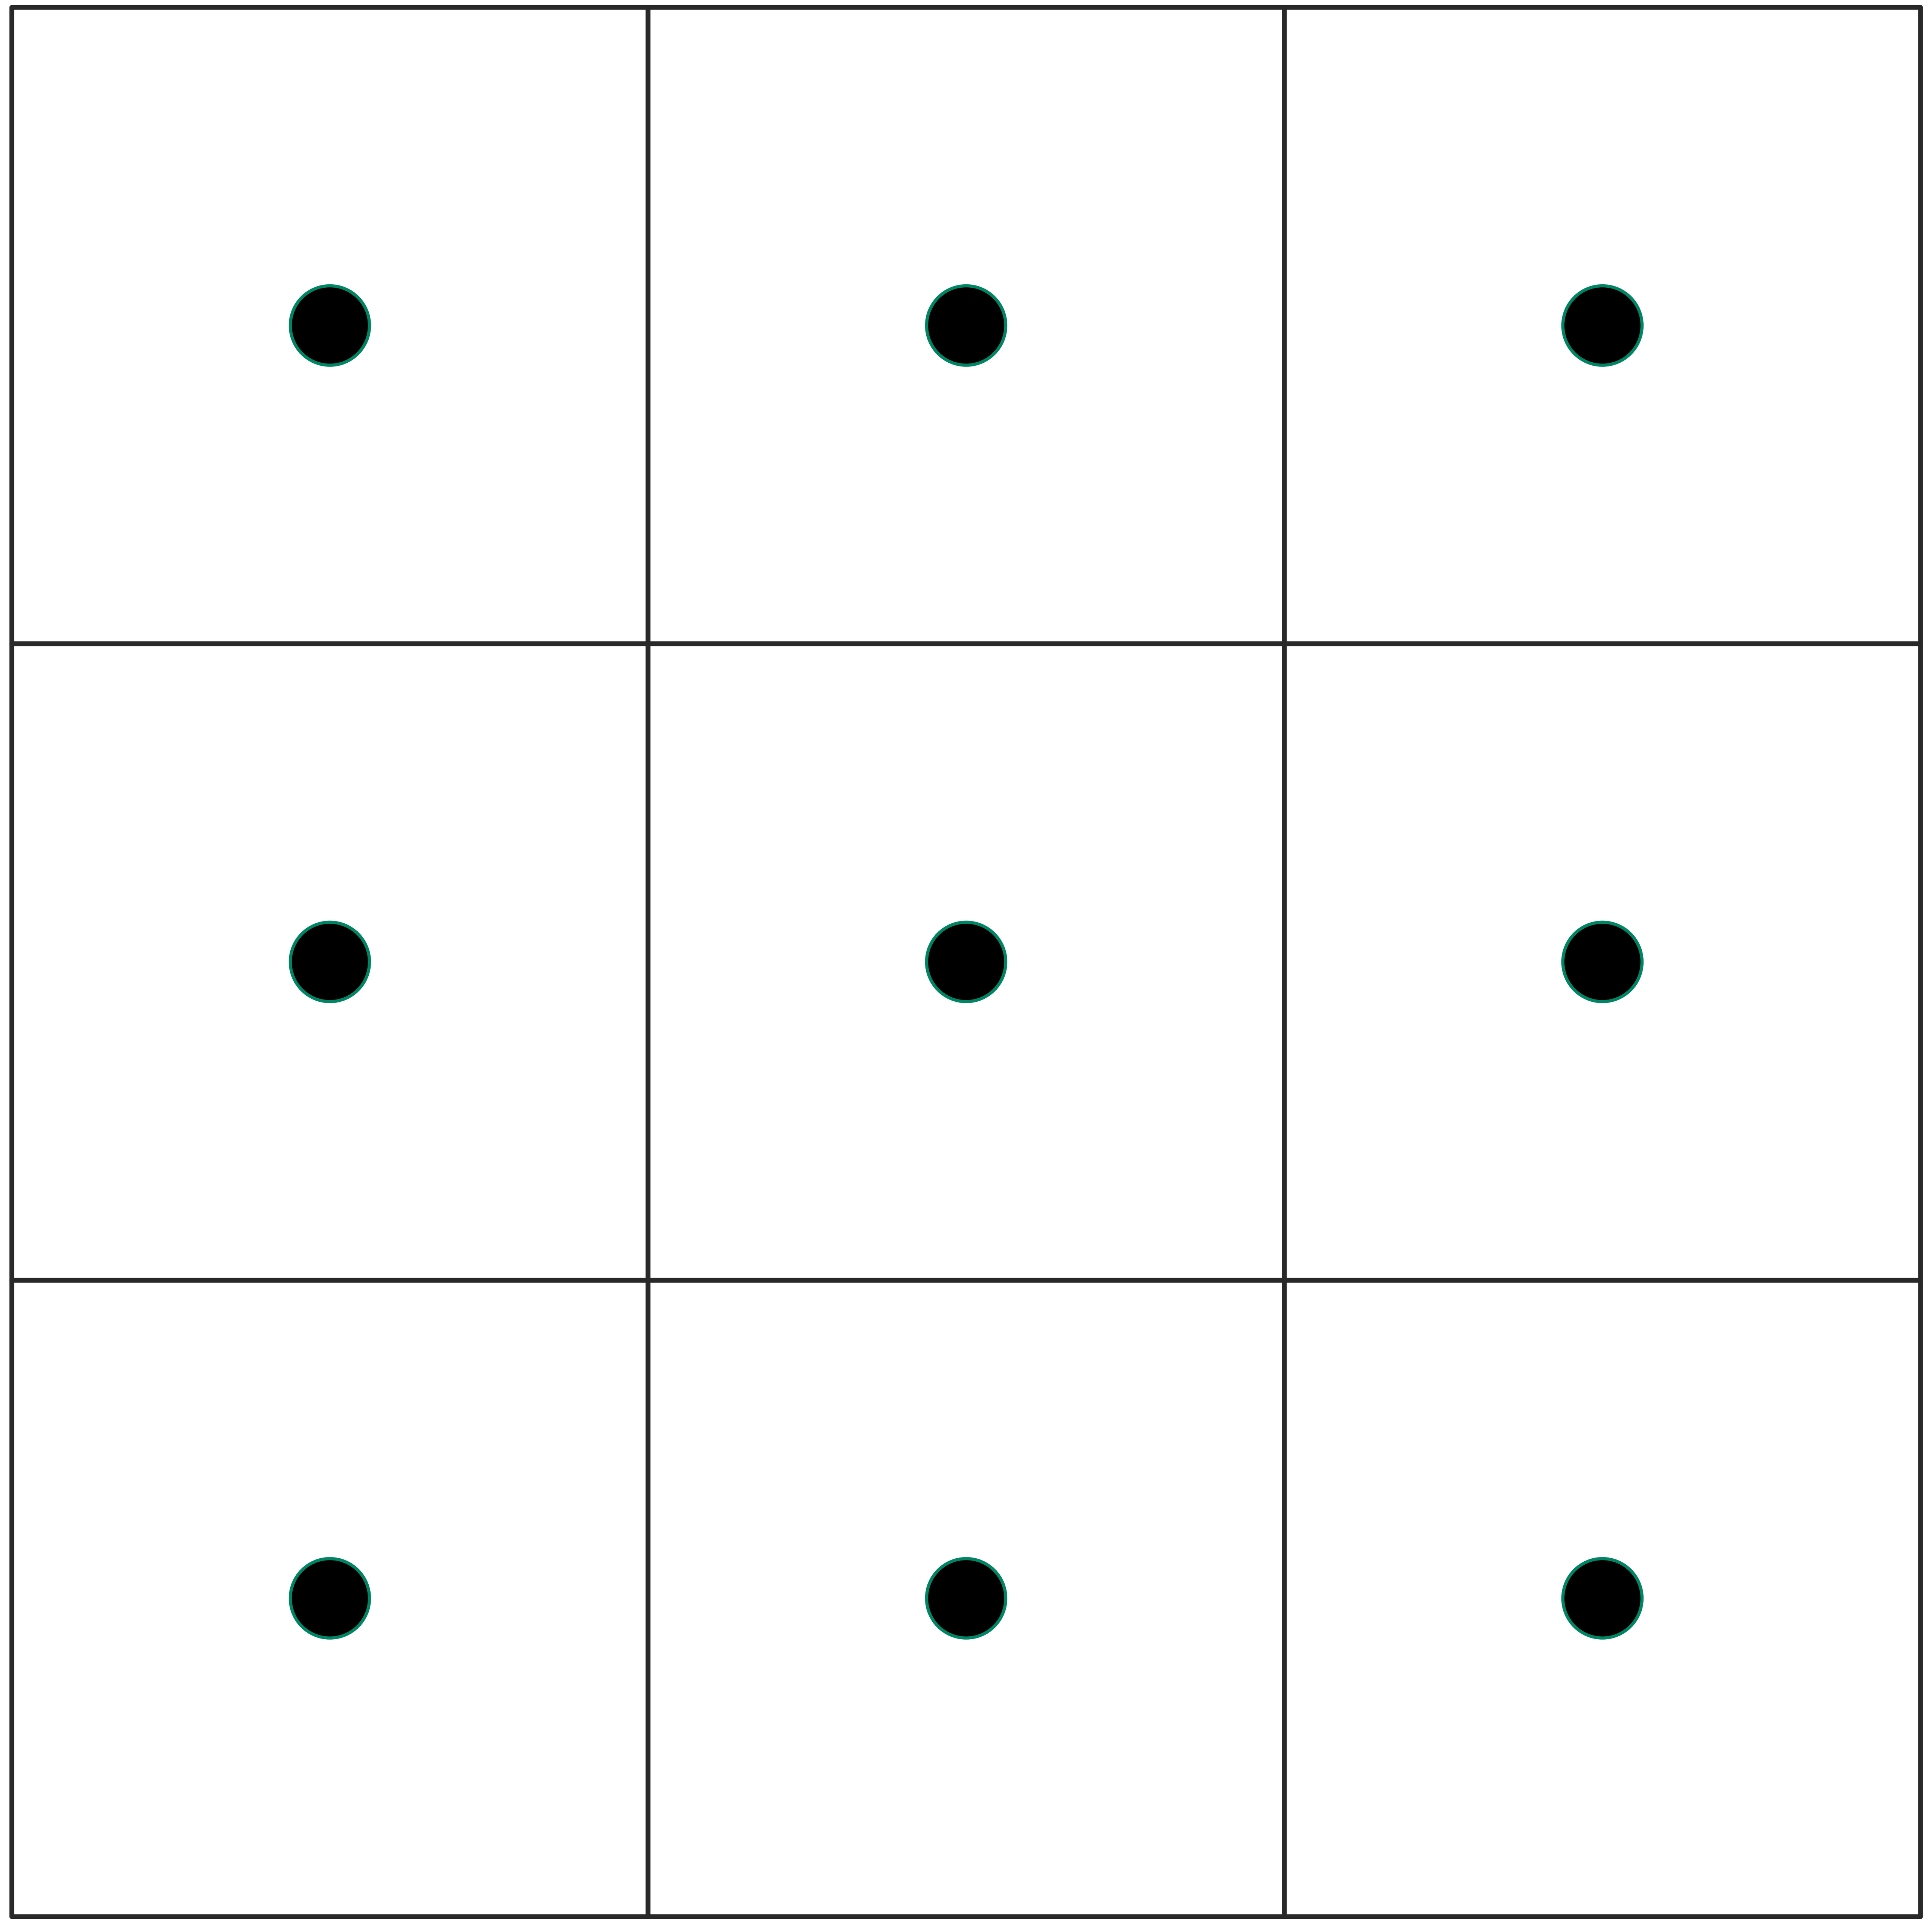

Texture mapping – why we need filtering?

- Texture – a 1D/2D/3D/4D grid of discrete values
- Values defined only at **texel centers**
- "Pixel/texel is not a little square!"
- Infinitely small point – **Dirac delta**

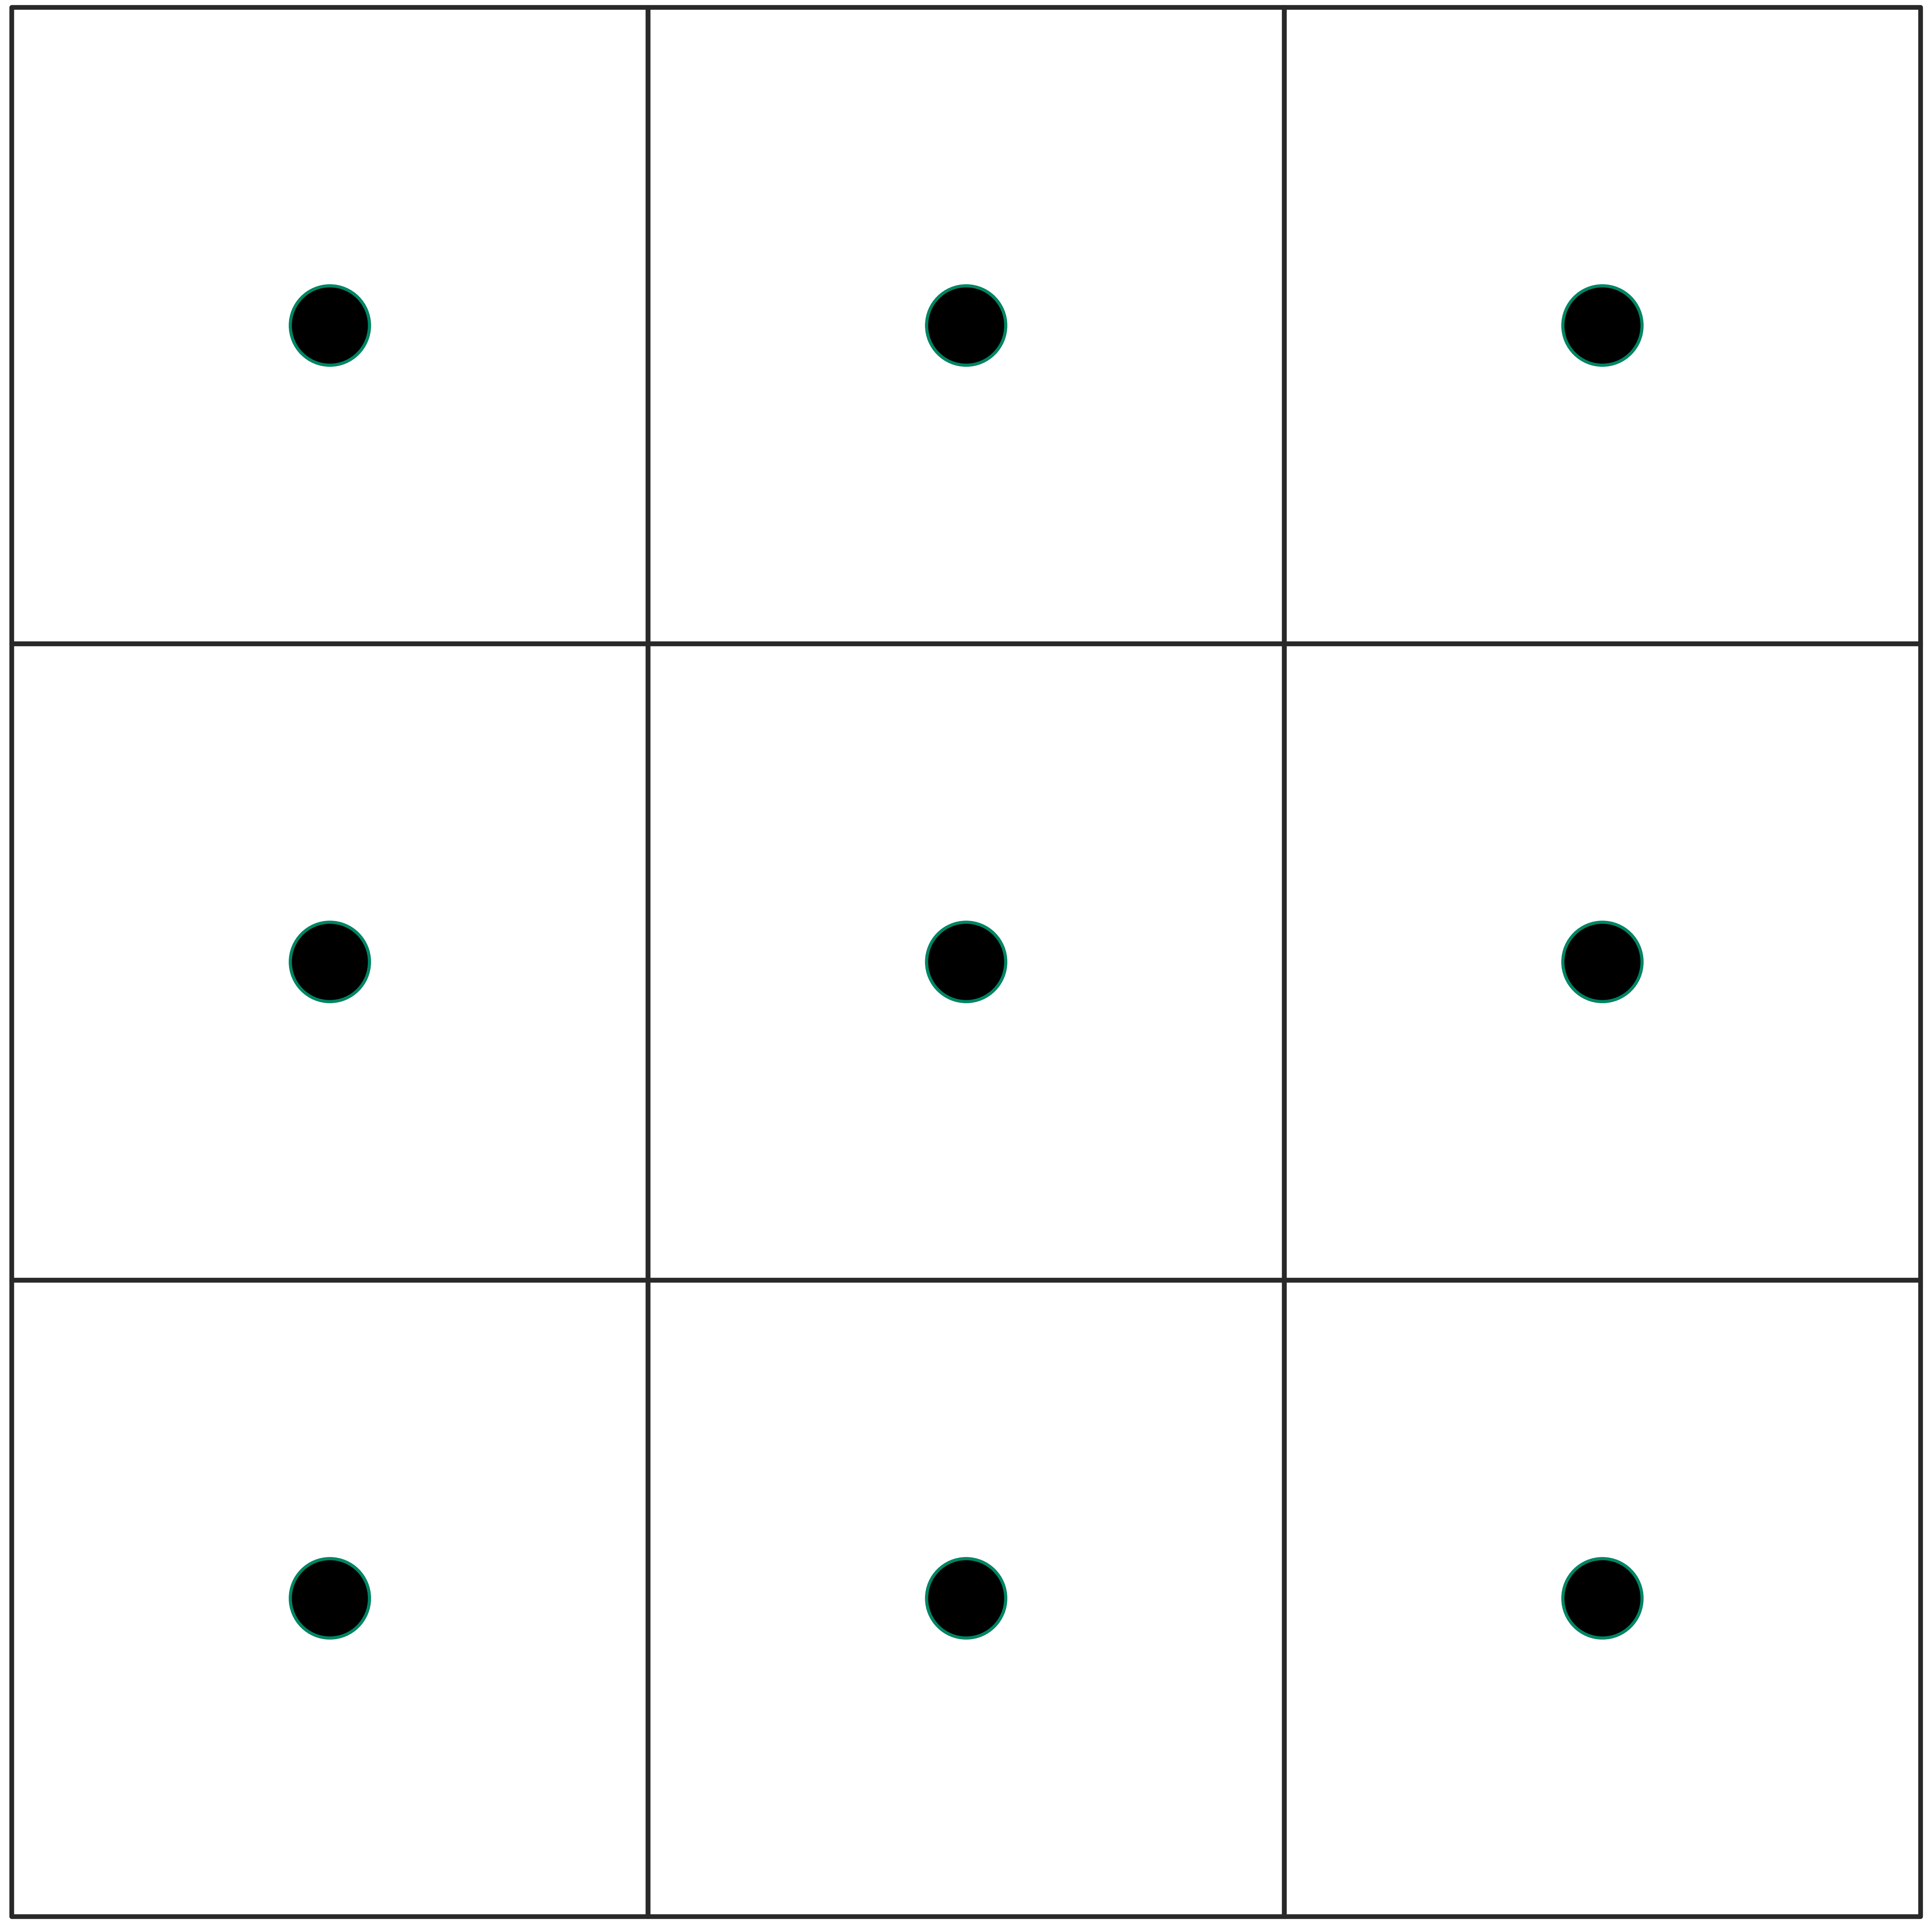

Texture mapping – why we need filtering?

- Texture – a 1D/2D/3D/4D grid of discrete values
- Values defined only at **texel centers**
- "Pixel/texel is not a little square!"
- Infinitely small point – **Dirac delta**
- What happens between?

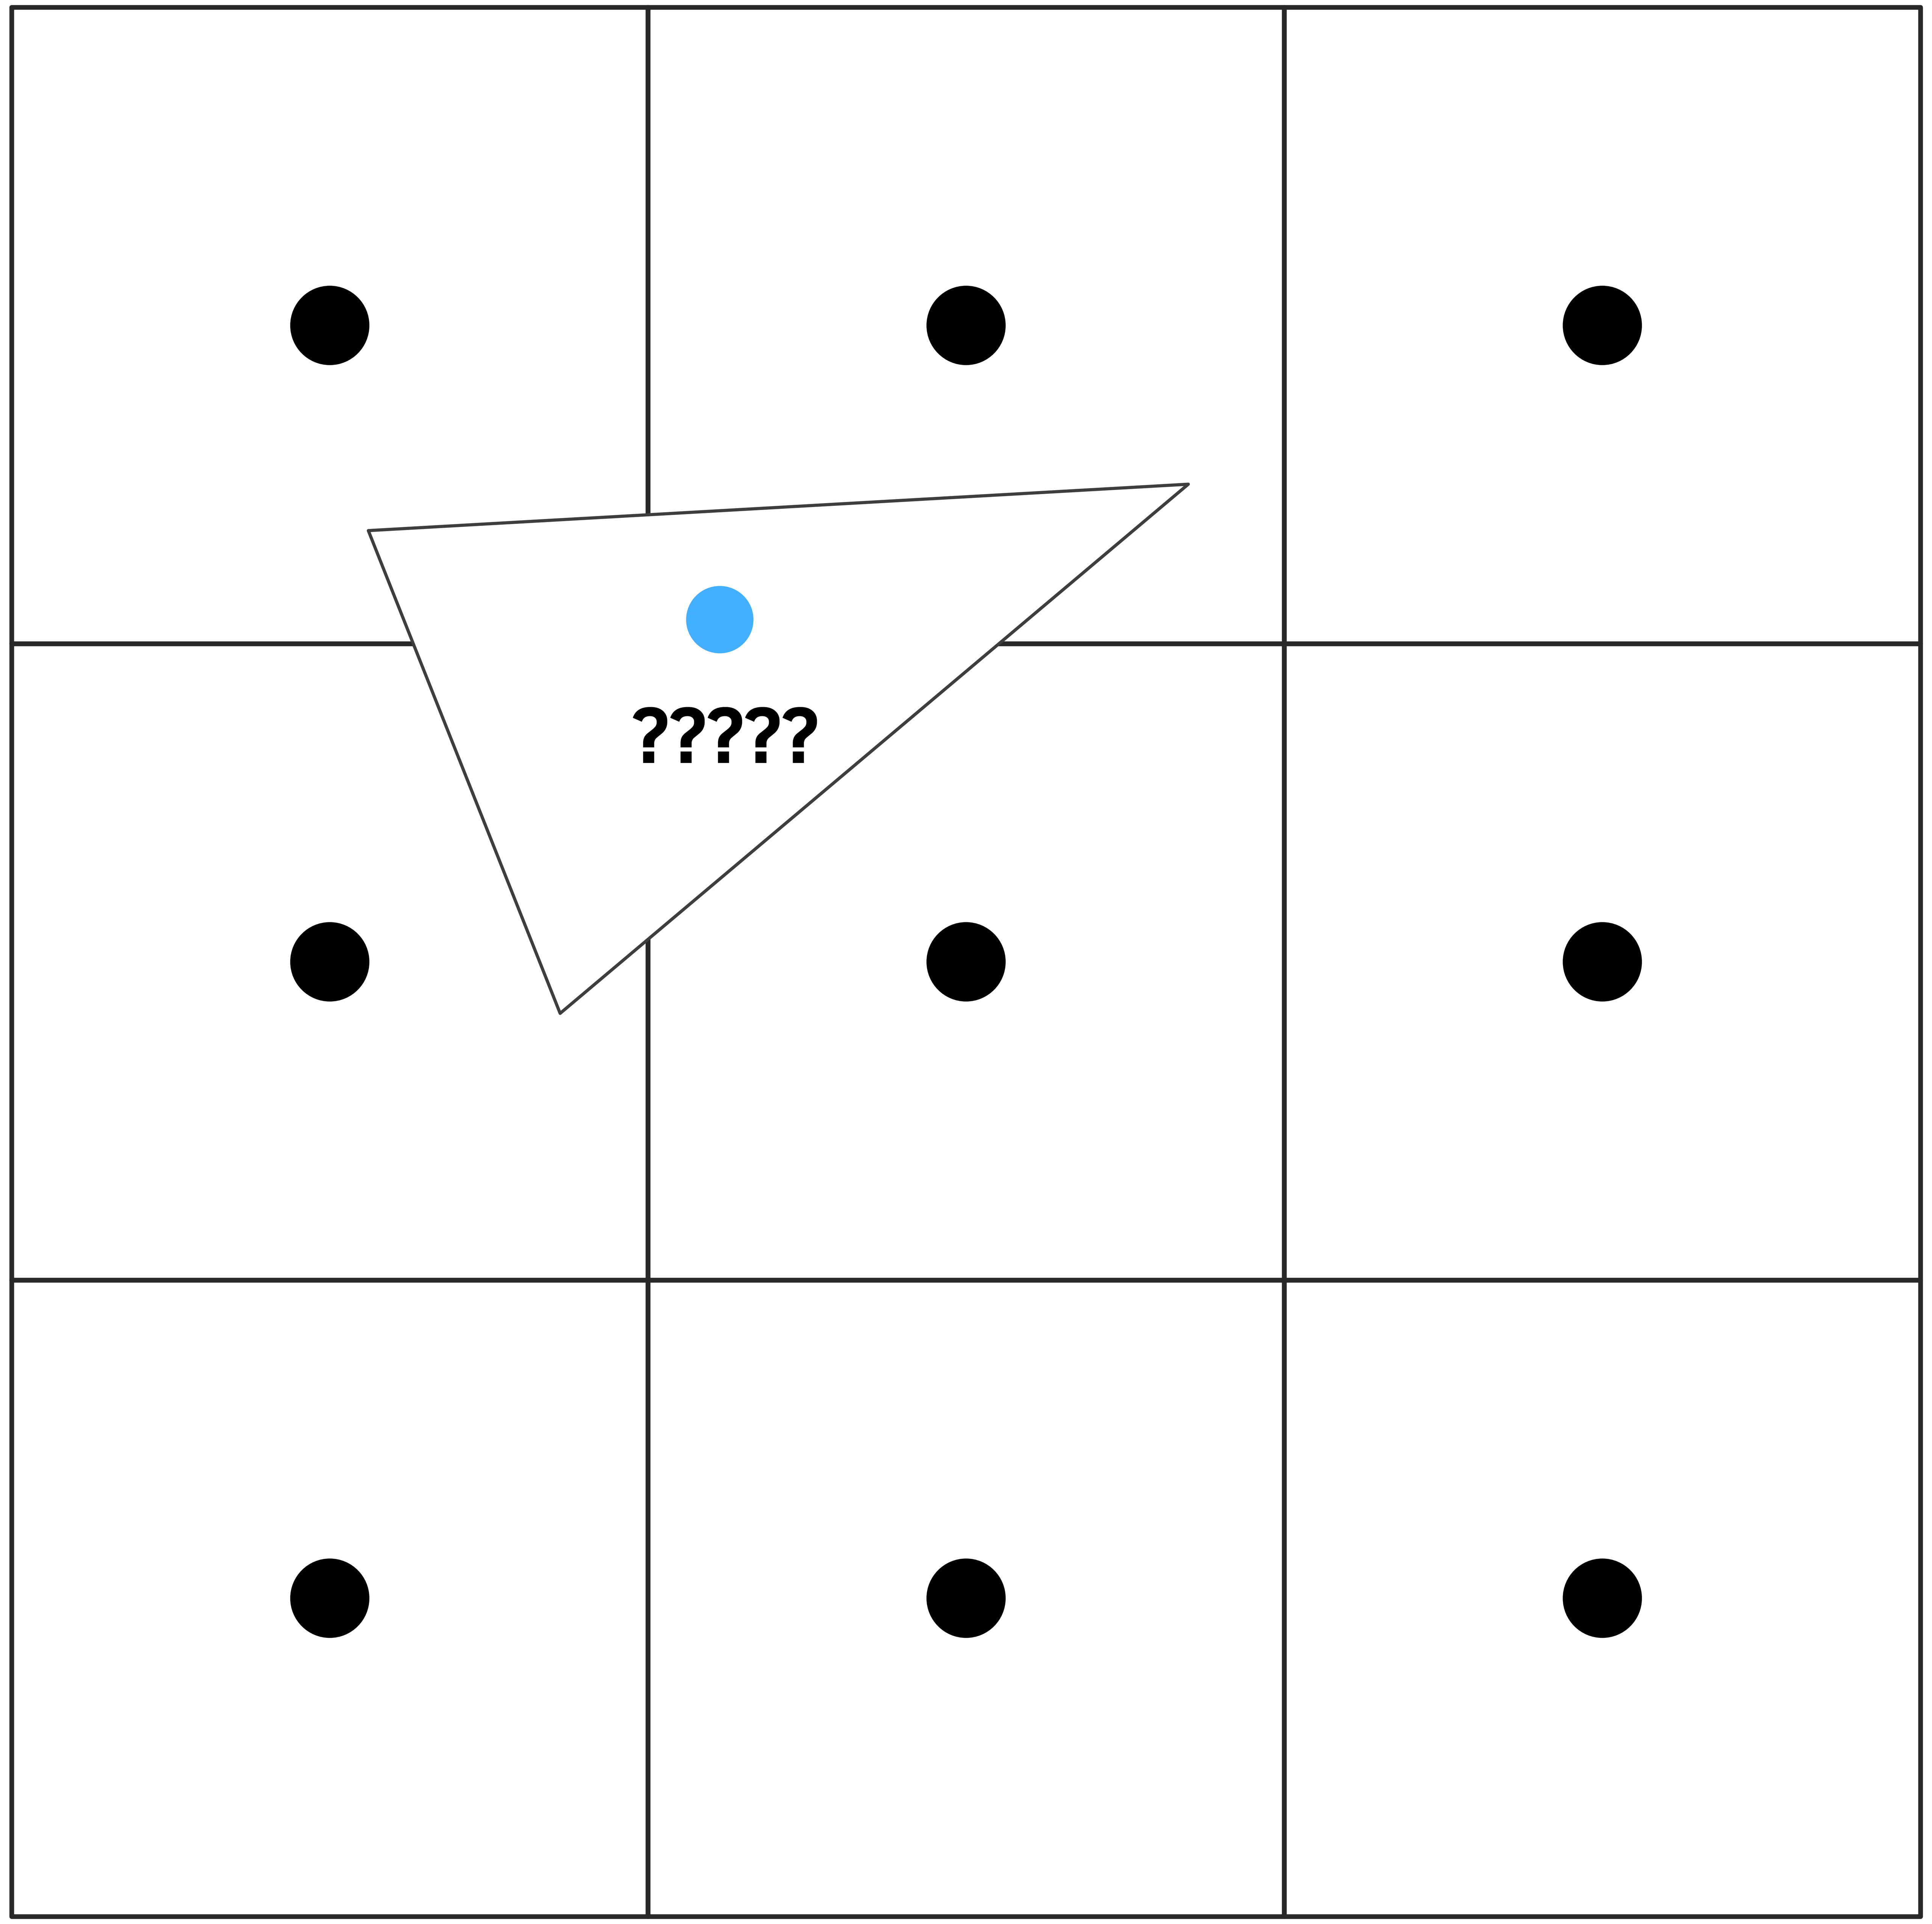

Texture mapping – why we need filtering?

- Texture – a 1D/2D/3D/4D grid of discrete values
- Values defined only at **texel centers**
- “Pixel/texel is not a little square!”
- Infinitely small point – **Dirac delta**
- What happens between?
- **Filtering and interpolation**
- Weighted averaging of multiple texture samples

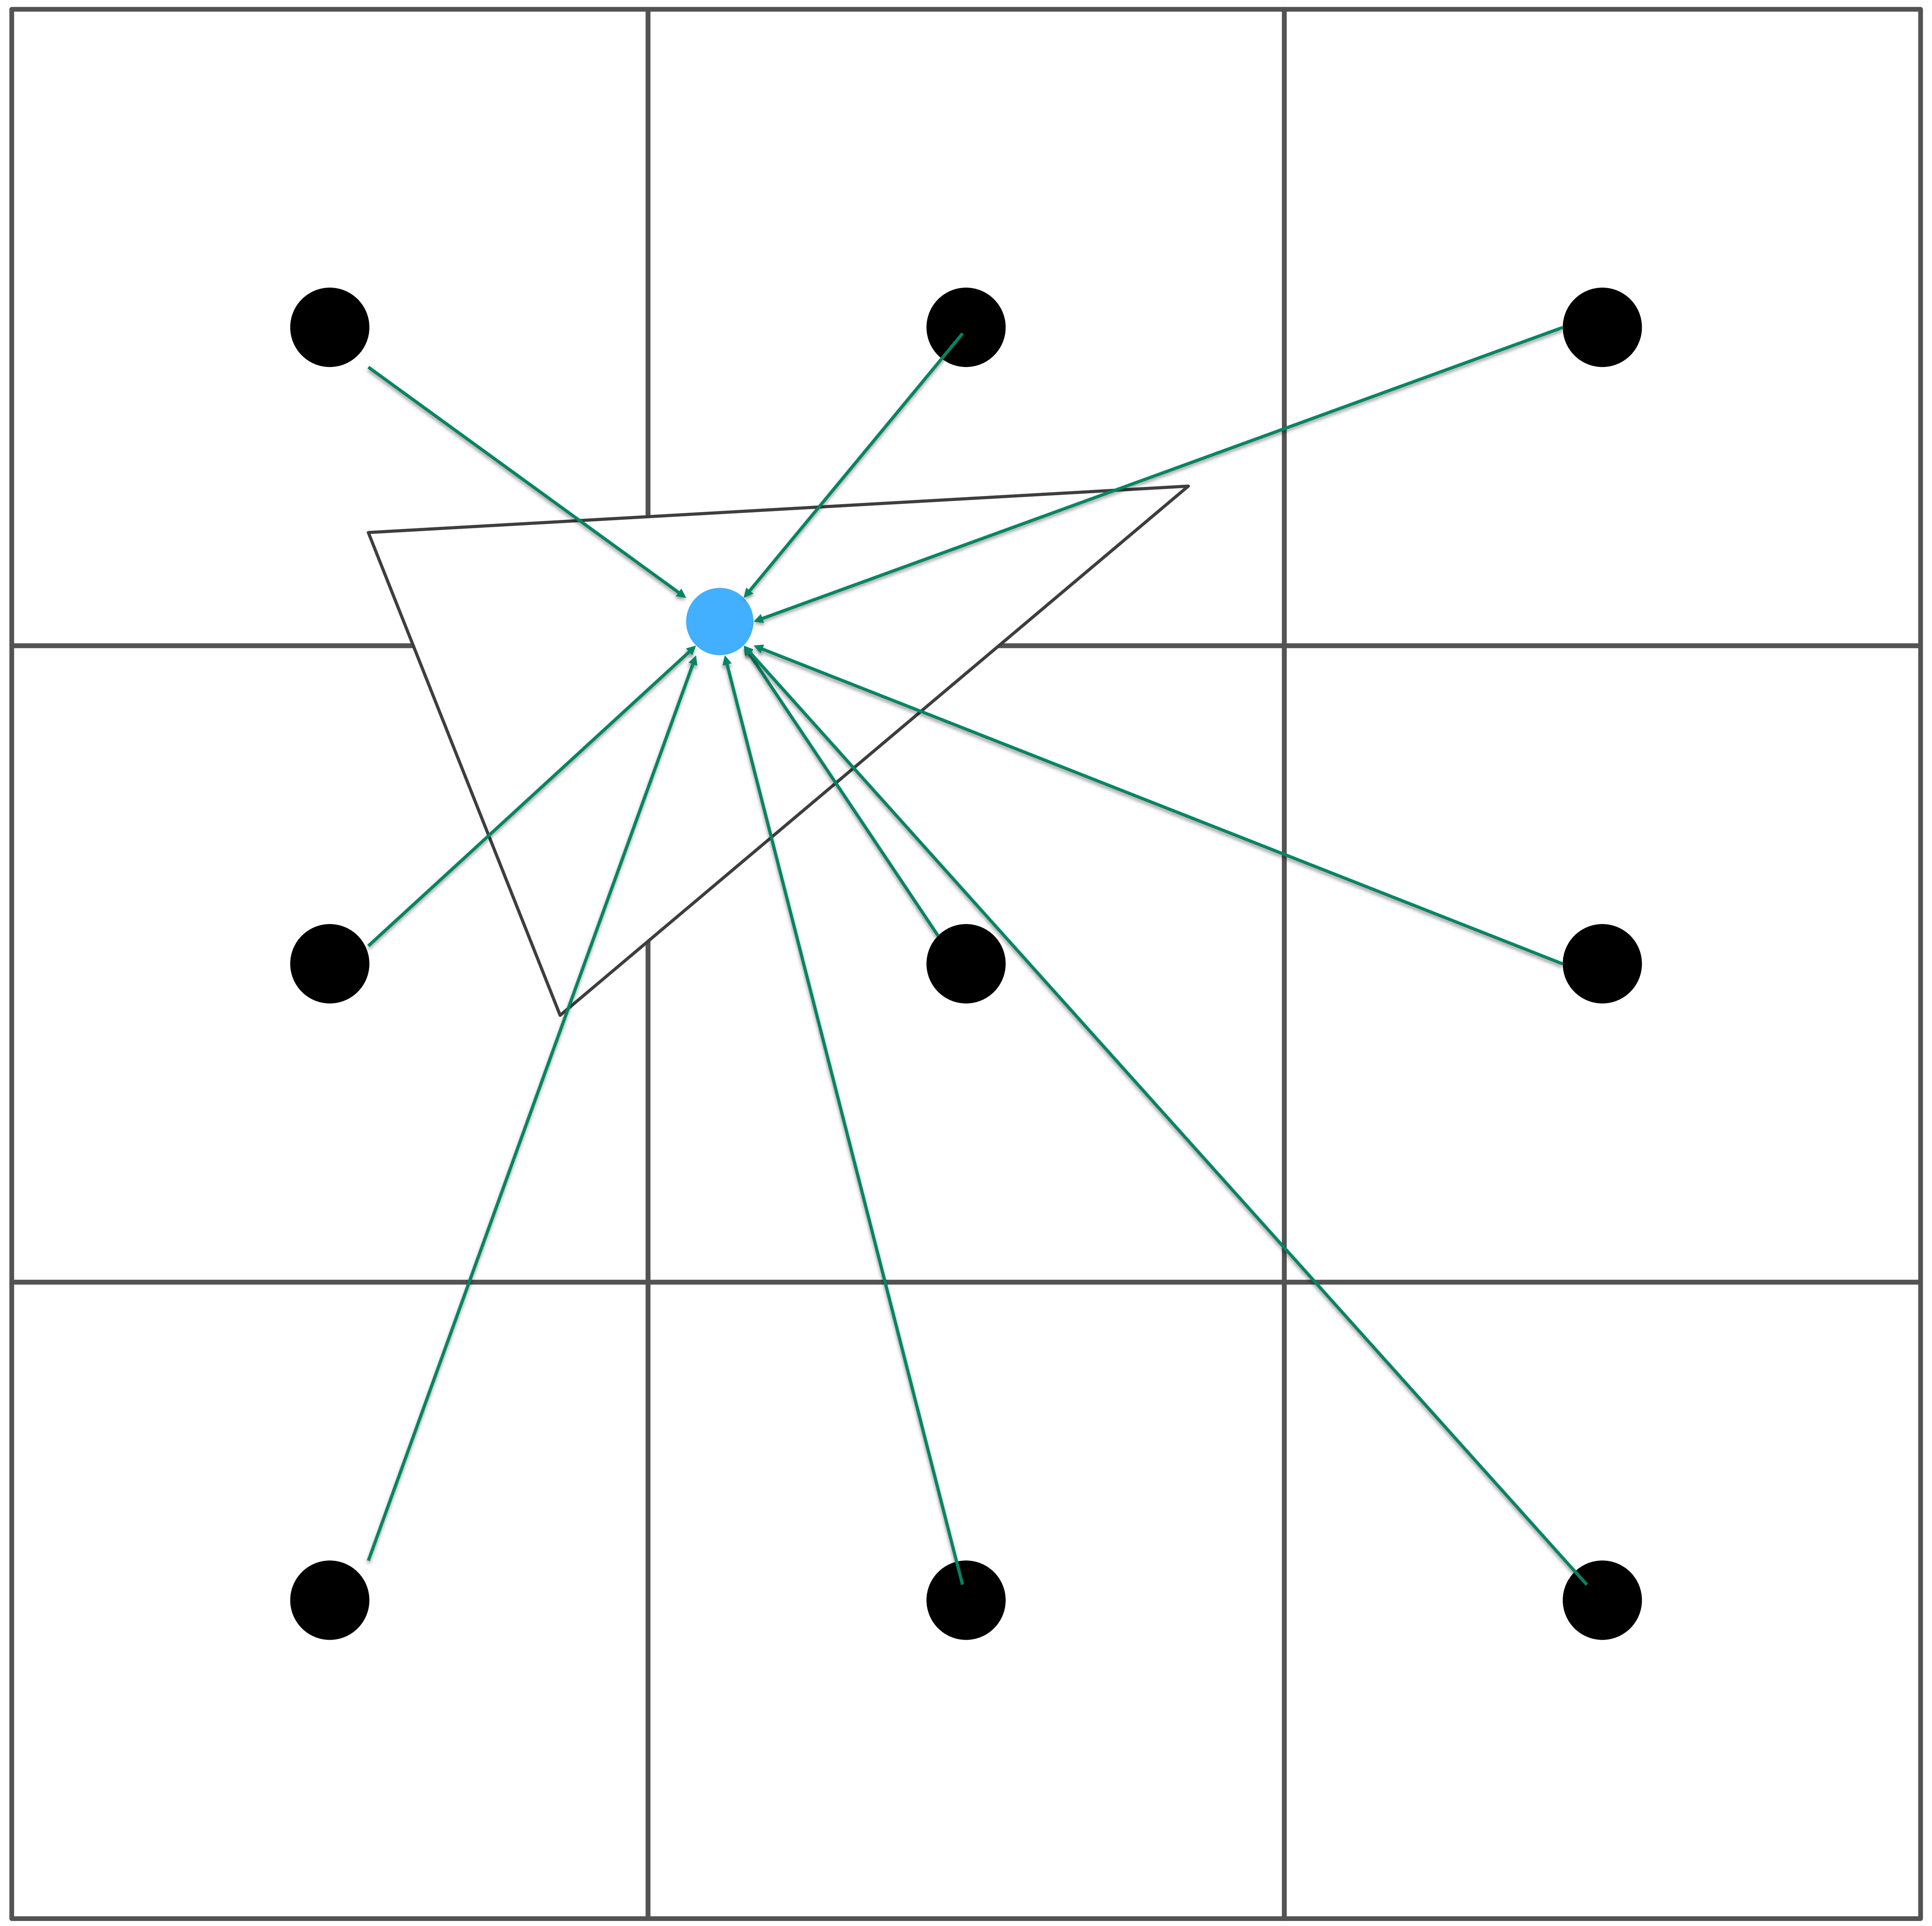

Texture filtering – minification

- Multiple texels might cover a single pixel area
- Potentially thousands (millions?) texels

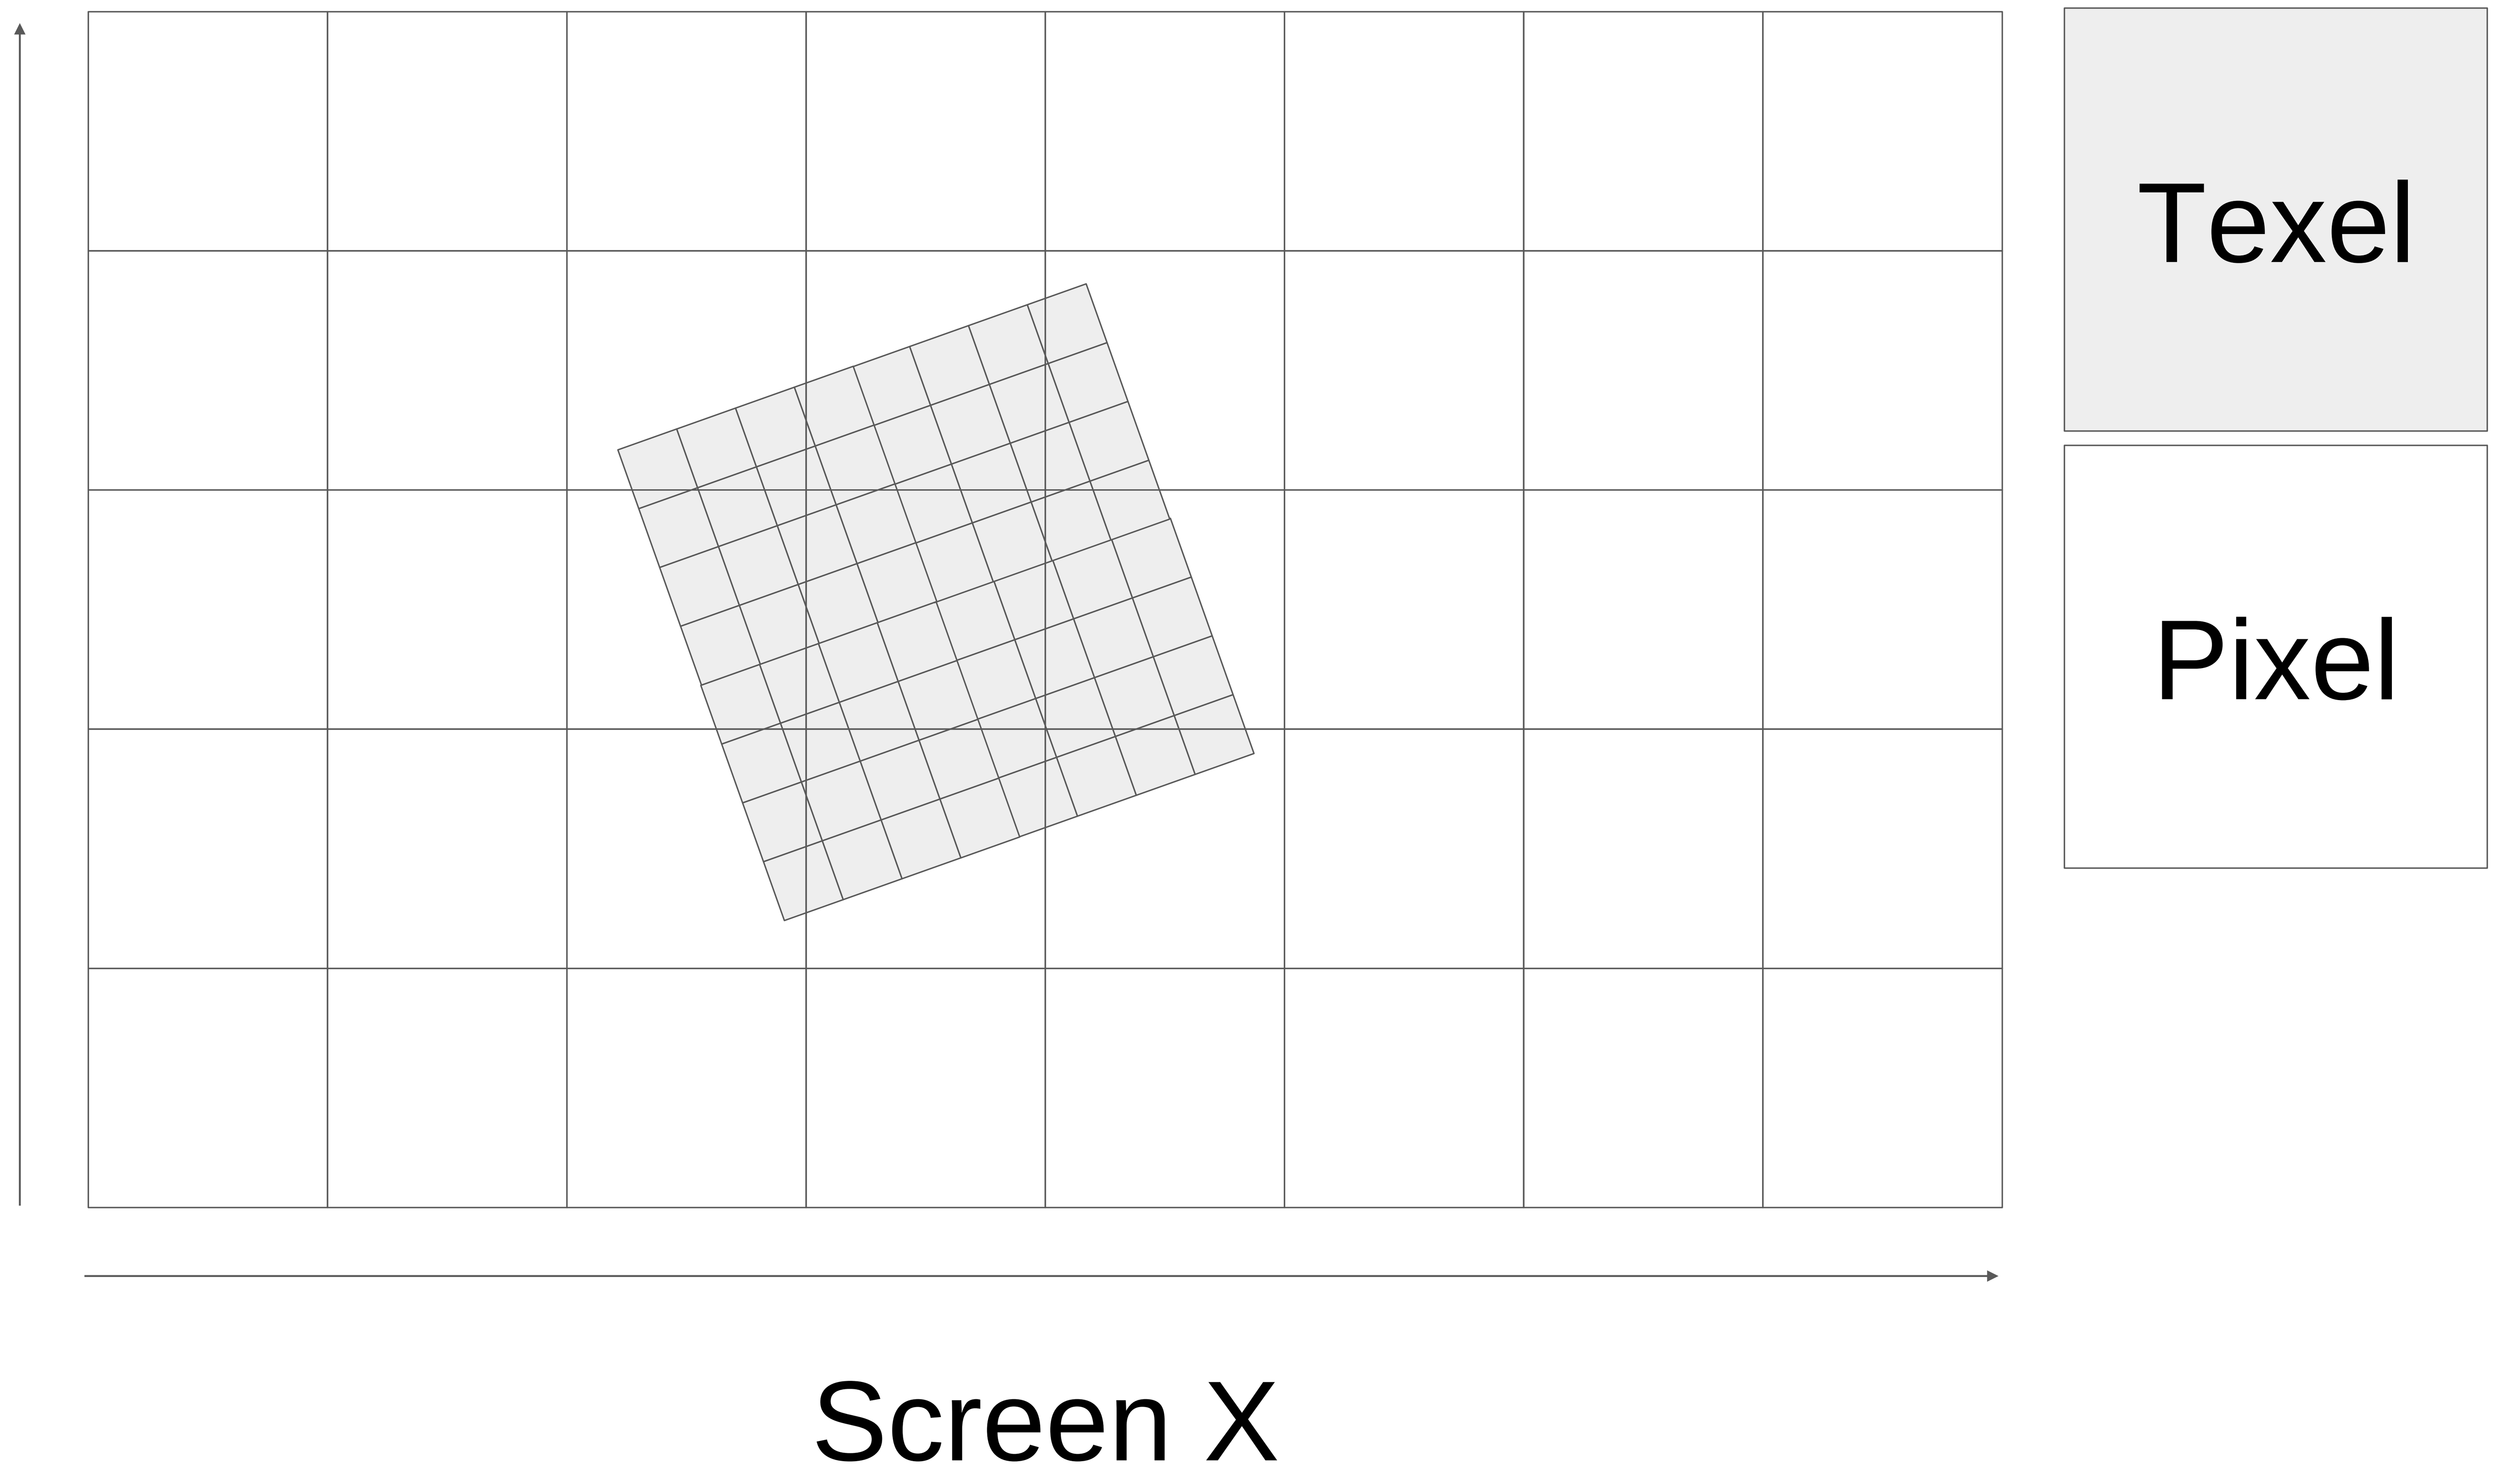

Texture filtering – minification

- Multiple texels might cover a single pixel area
- Potentially thousands (millions?) texels

Solution – **full filter** – EWA, anisotropic filtering

- Possibly very slow
- Hundreds+ of texture samples
- Higher quality

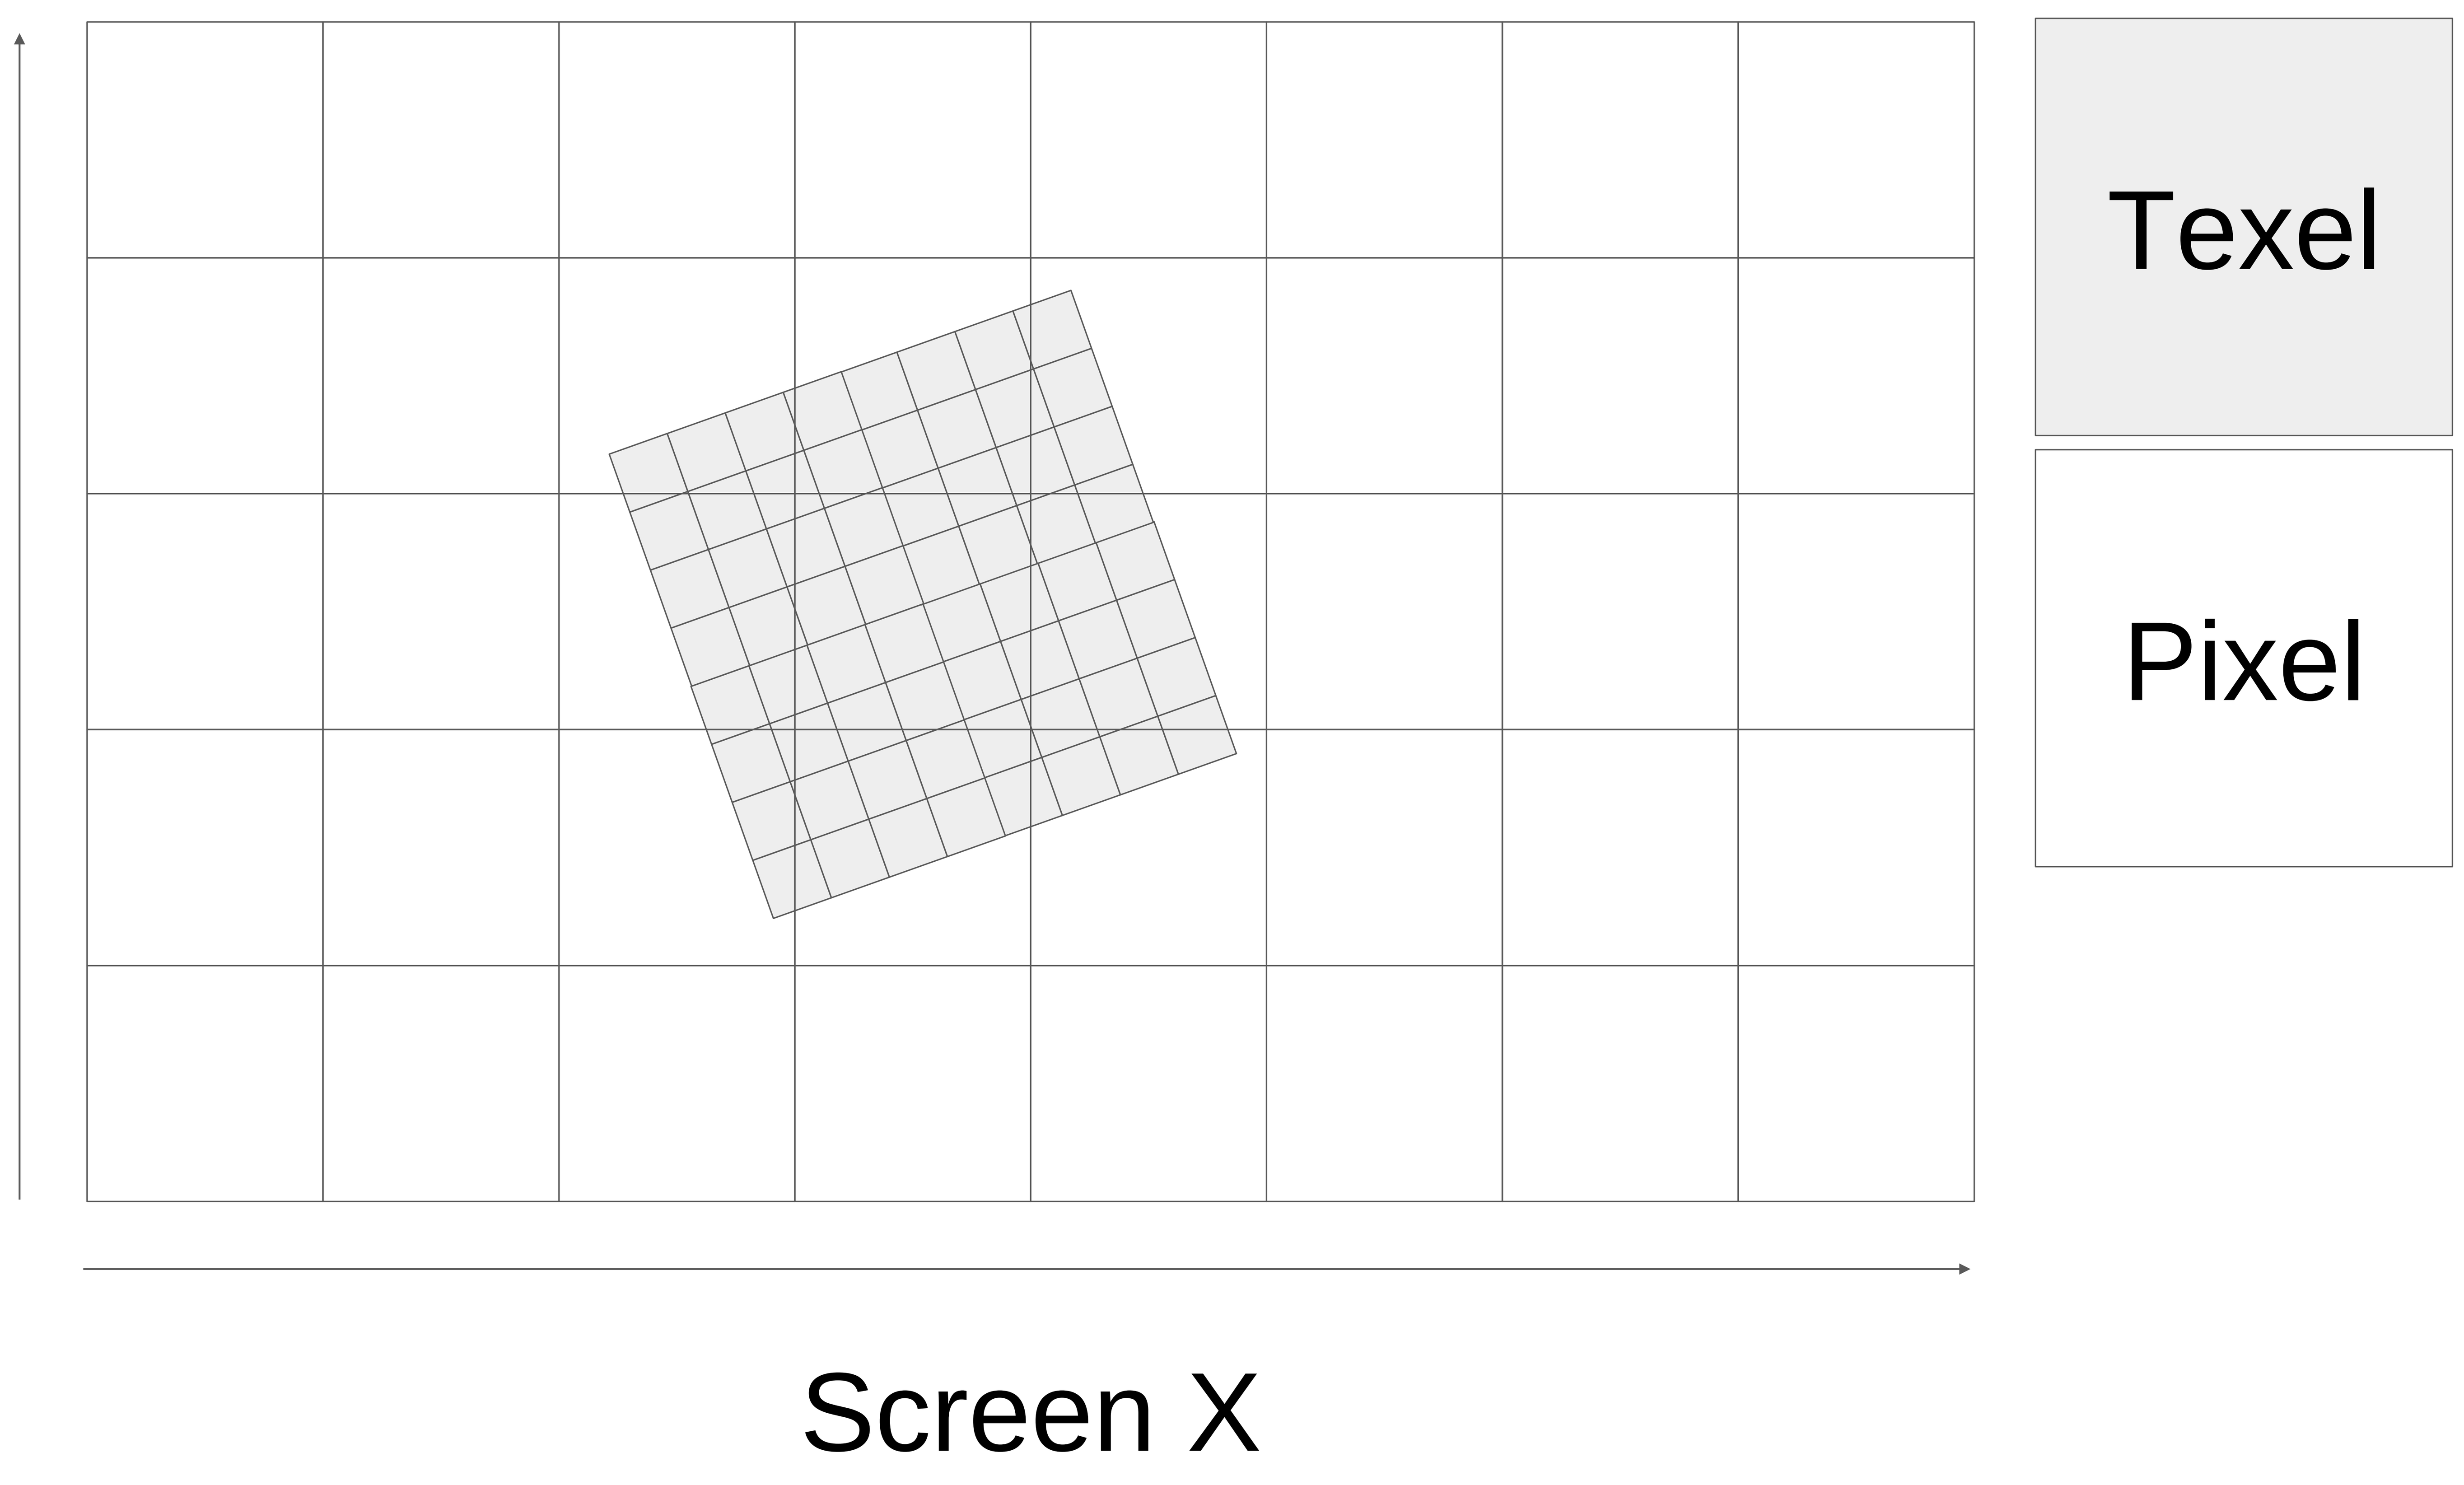

Texture filtering – minification

- Multiple texels might cover a single pixel area
- Potentially thousands (millions?) texels

Solution – **full filter** – EWA, anisotropic filtering

- Possibly very slow
- Hundreds+ of texture samples
- Higher quality

Solution – **prefilter** – precomputed mipmap pyramid

- Very fast
- Low quality (blurry!)

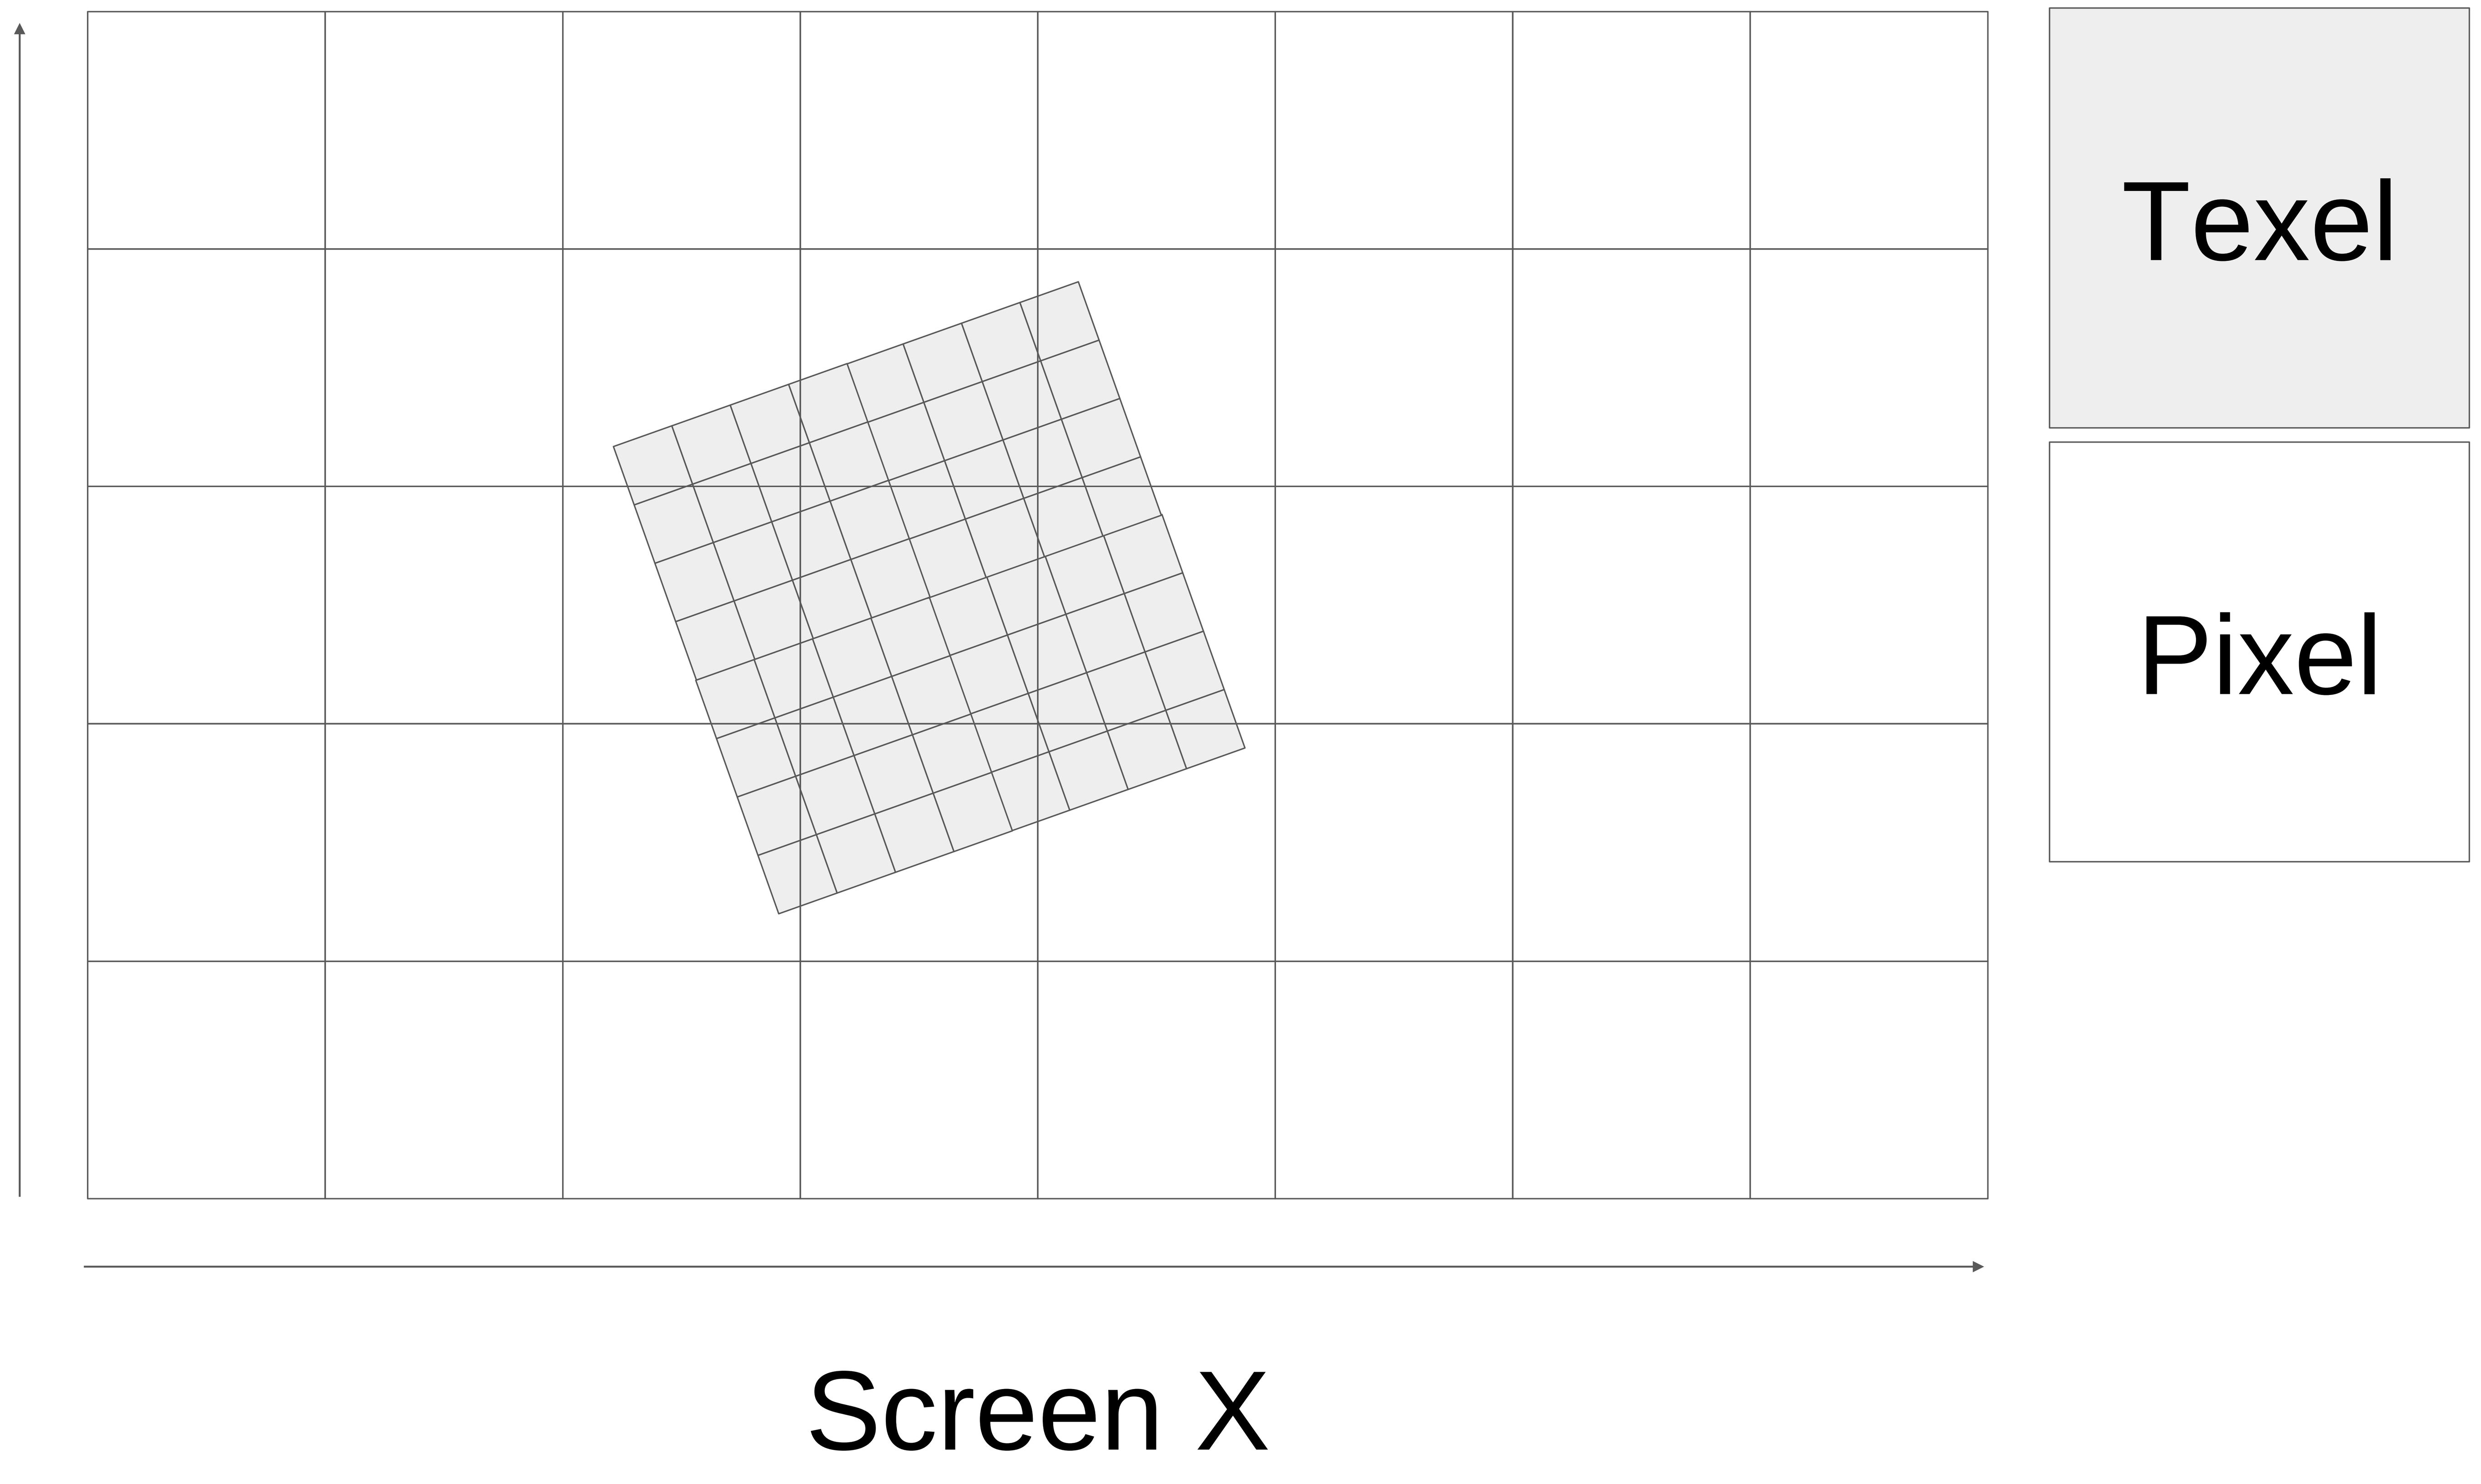

Texture filtering – minification

- Multiple texels might cover a single pixel area
- Potentially thousands (millions?) texels

Solution – **full filter** – EWA, anisotropic filtering

- Possibly very slow
- Hundreds+ of texture samples
- Higher quality

Solution – **prefilter** – precomputed mipmap pyramid

- Very fast
- Low quality (blurry!)

(Often – hybrid, mipmapping + anisotropic/mip bias)

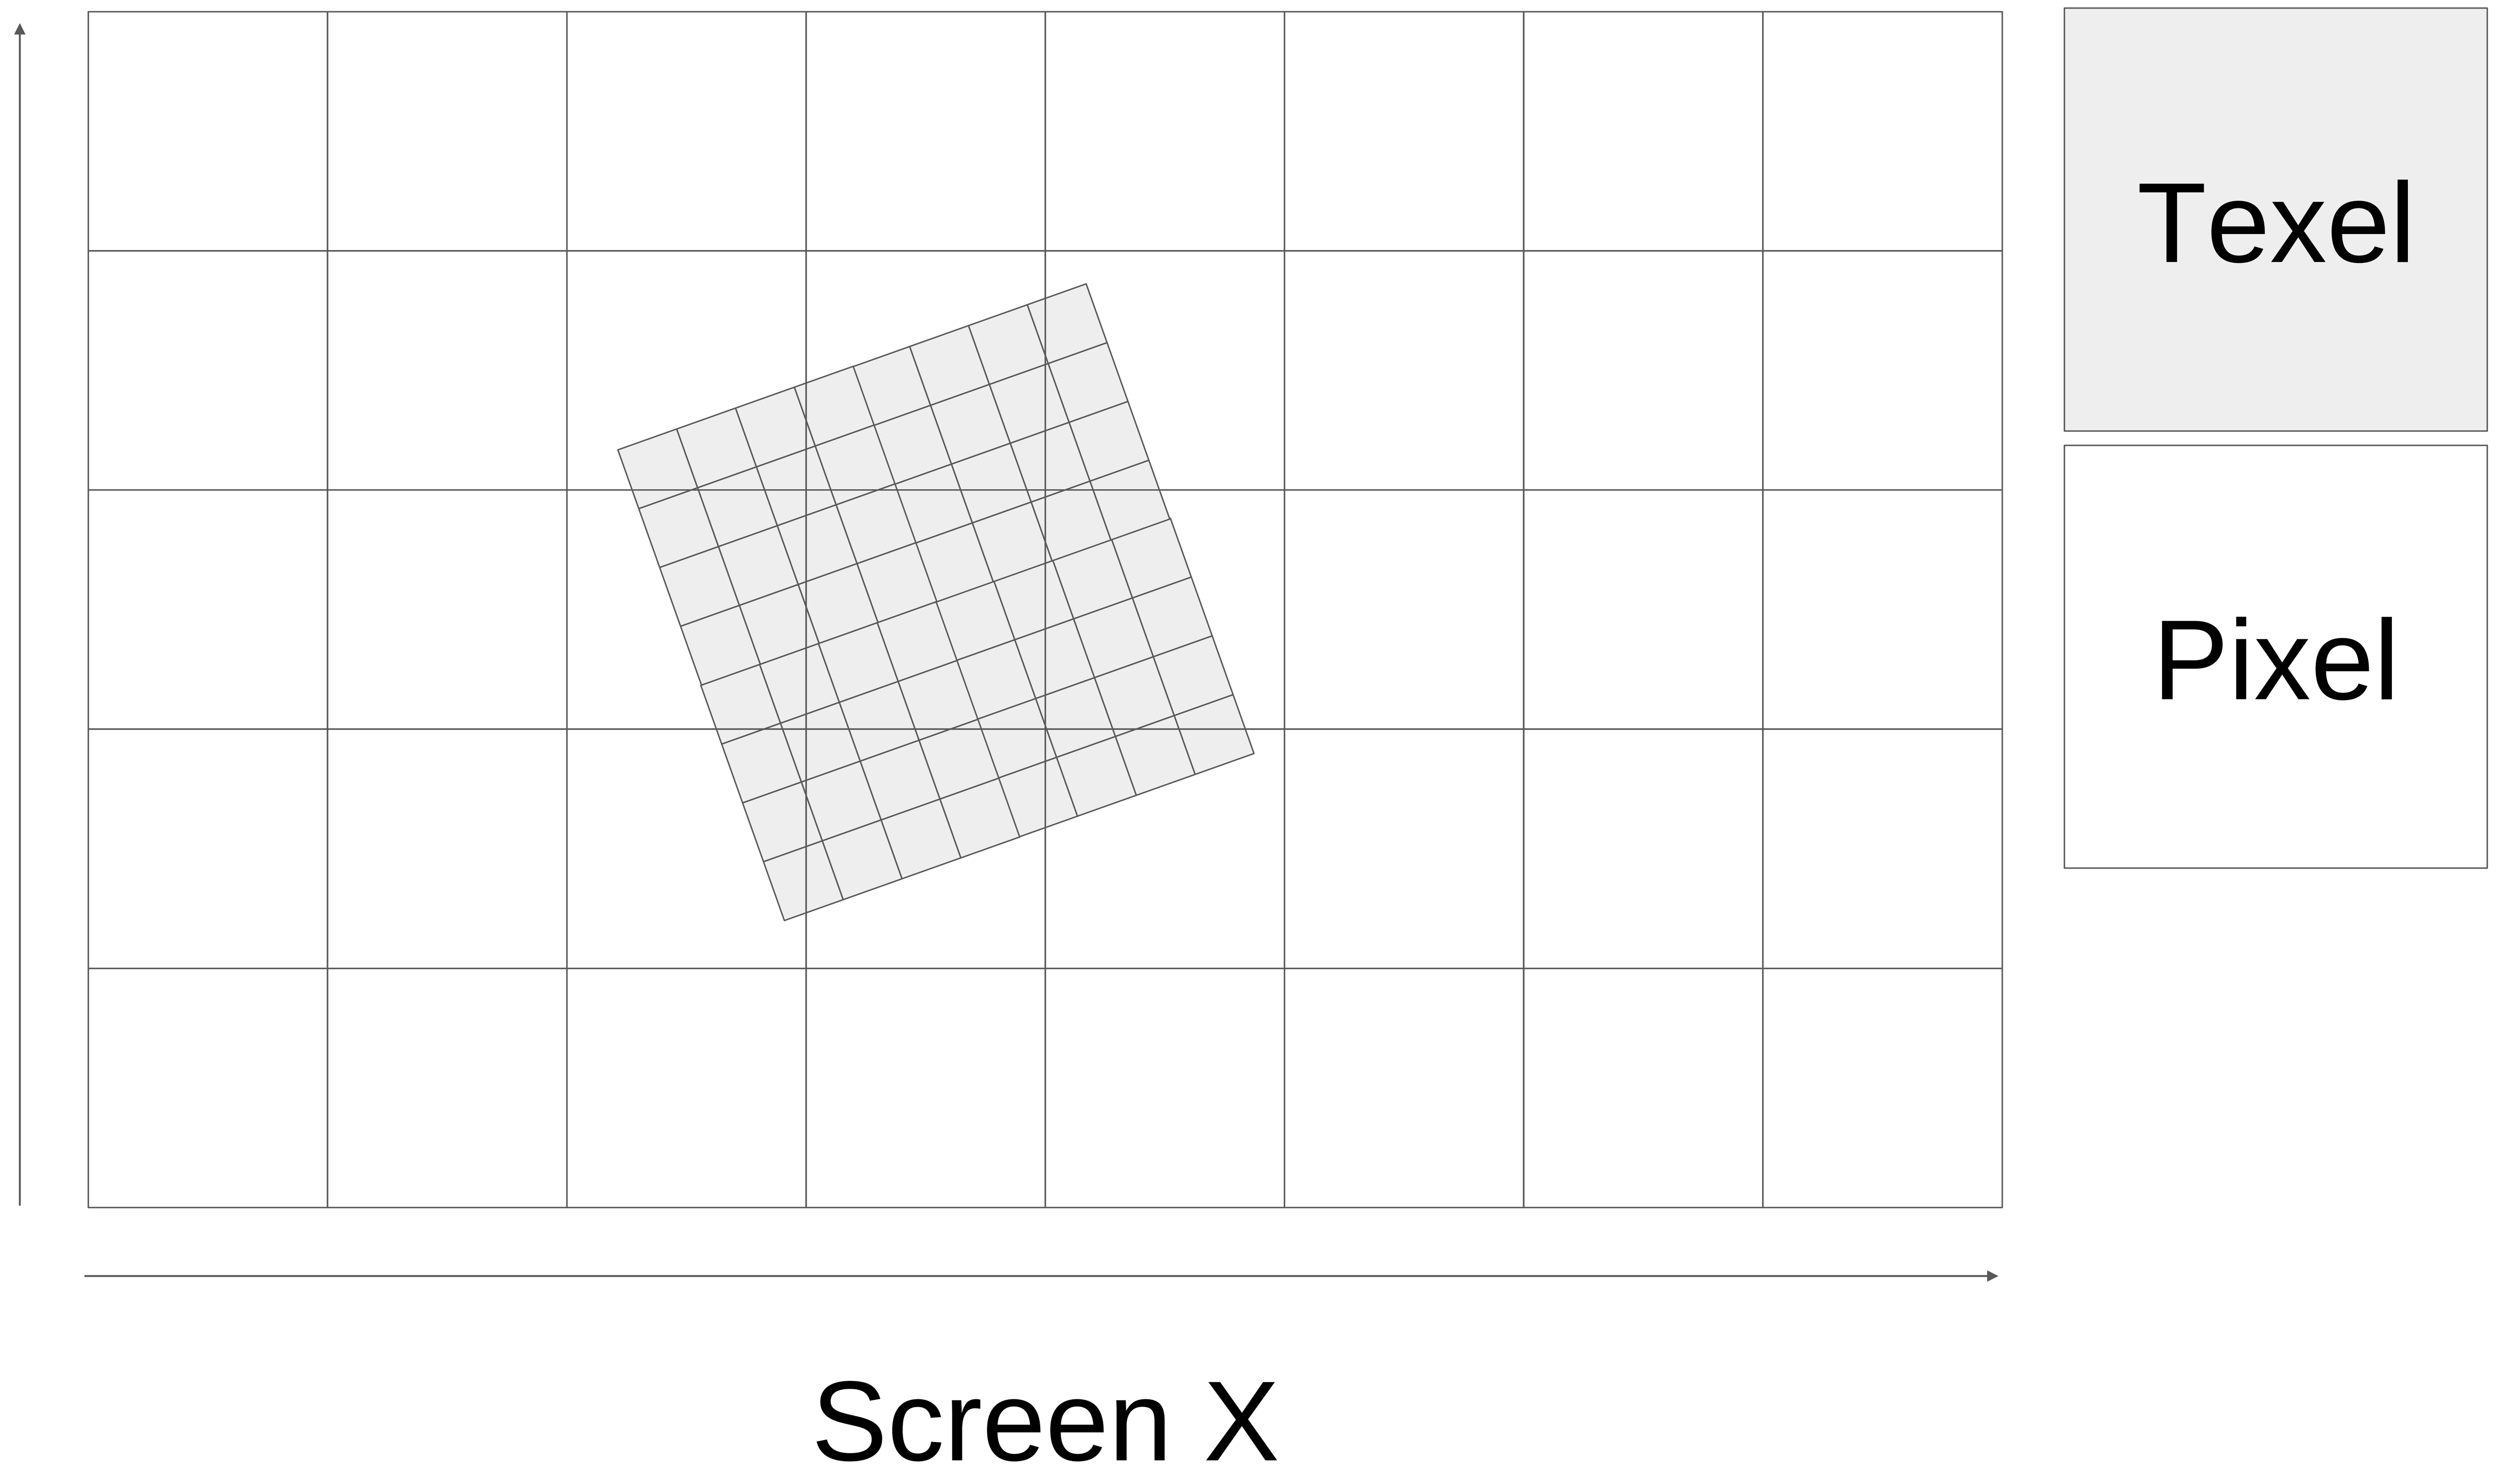

“Common knowledge”

Almost “axiomatic”

All modern graphics APIs standardize filtering

- Standard filters – (low-quality) bi/trilinear, anisotropic

“Common knowledge”

Almost “axiomatic”

All modern graphics APIs standardize filtering

- Standard filters – (low-quality) bi/trilinear, anisotropic
- All modern GPUs have dedicated filtering hardware
- On most standardized formats -> zero cost for bi/trilinear
- Anisotropic still cheaper than naïve manual anisotropic

“Common knowledge”

Almost “axiomatic”

All modern graphics APIs standardize filtering

- Standard filters – (low-quality) bi/trilinear, anisotropic
- All modern GPUs have dedicated filtering hardware
- On most standardized formats -> zero cost for bi/trilinear
- Anisotropic still cheaper than naïve manual anisotropic
- Very easy and attractive to use without questioning

Feedback loop!

Motivation

Project beginning : Stochastic neural texture filtering (performance)

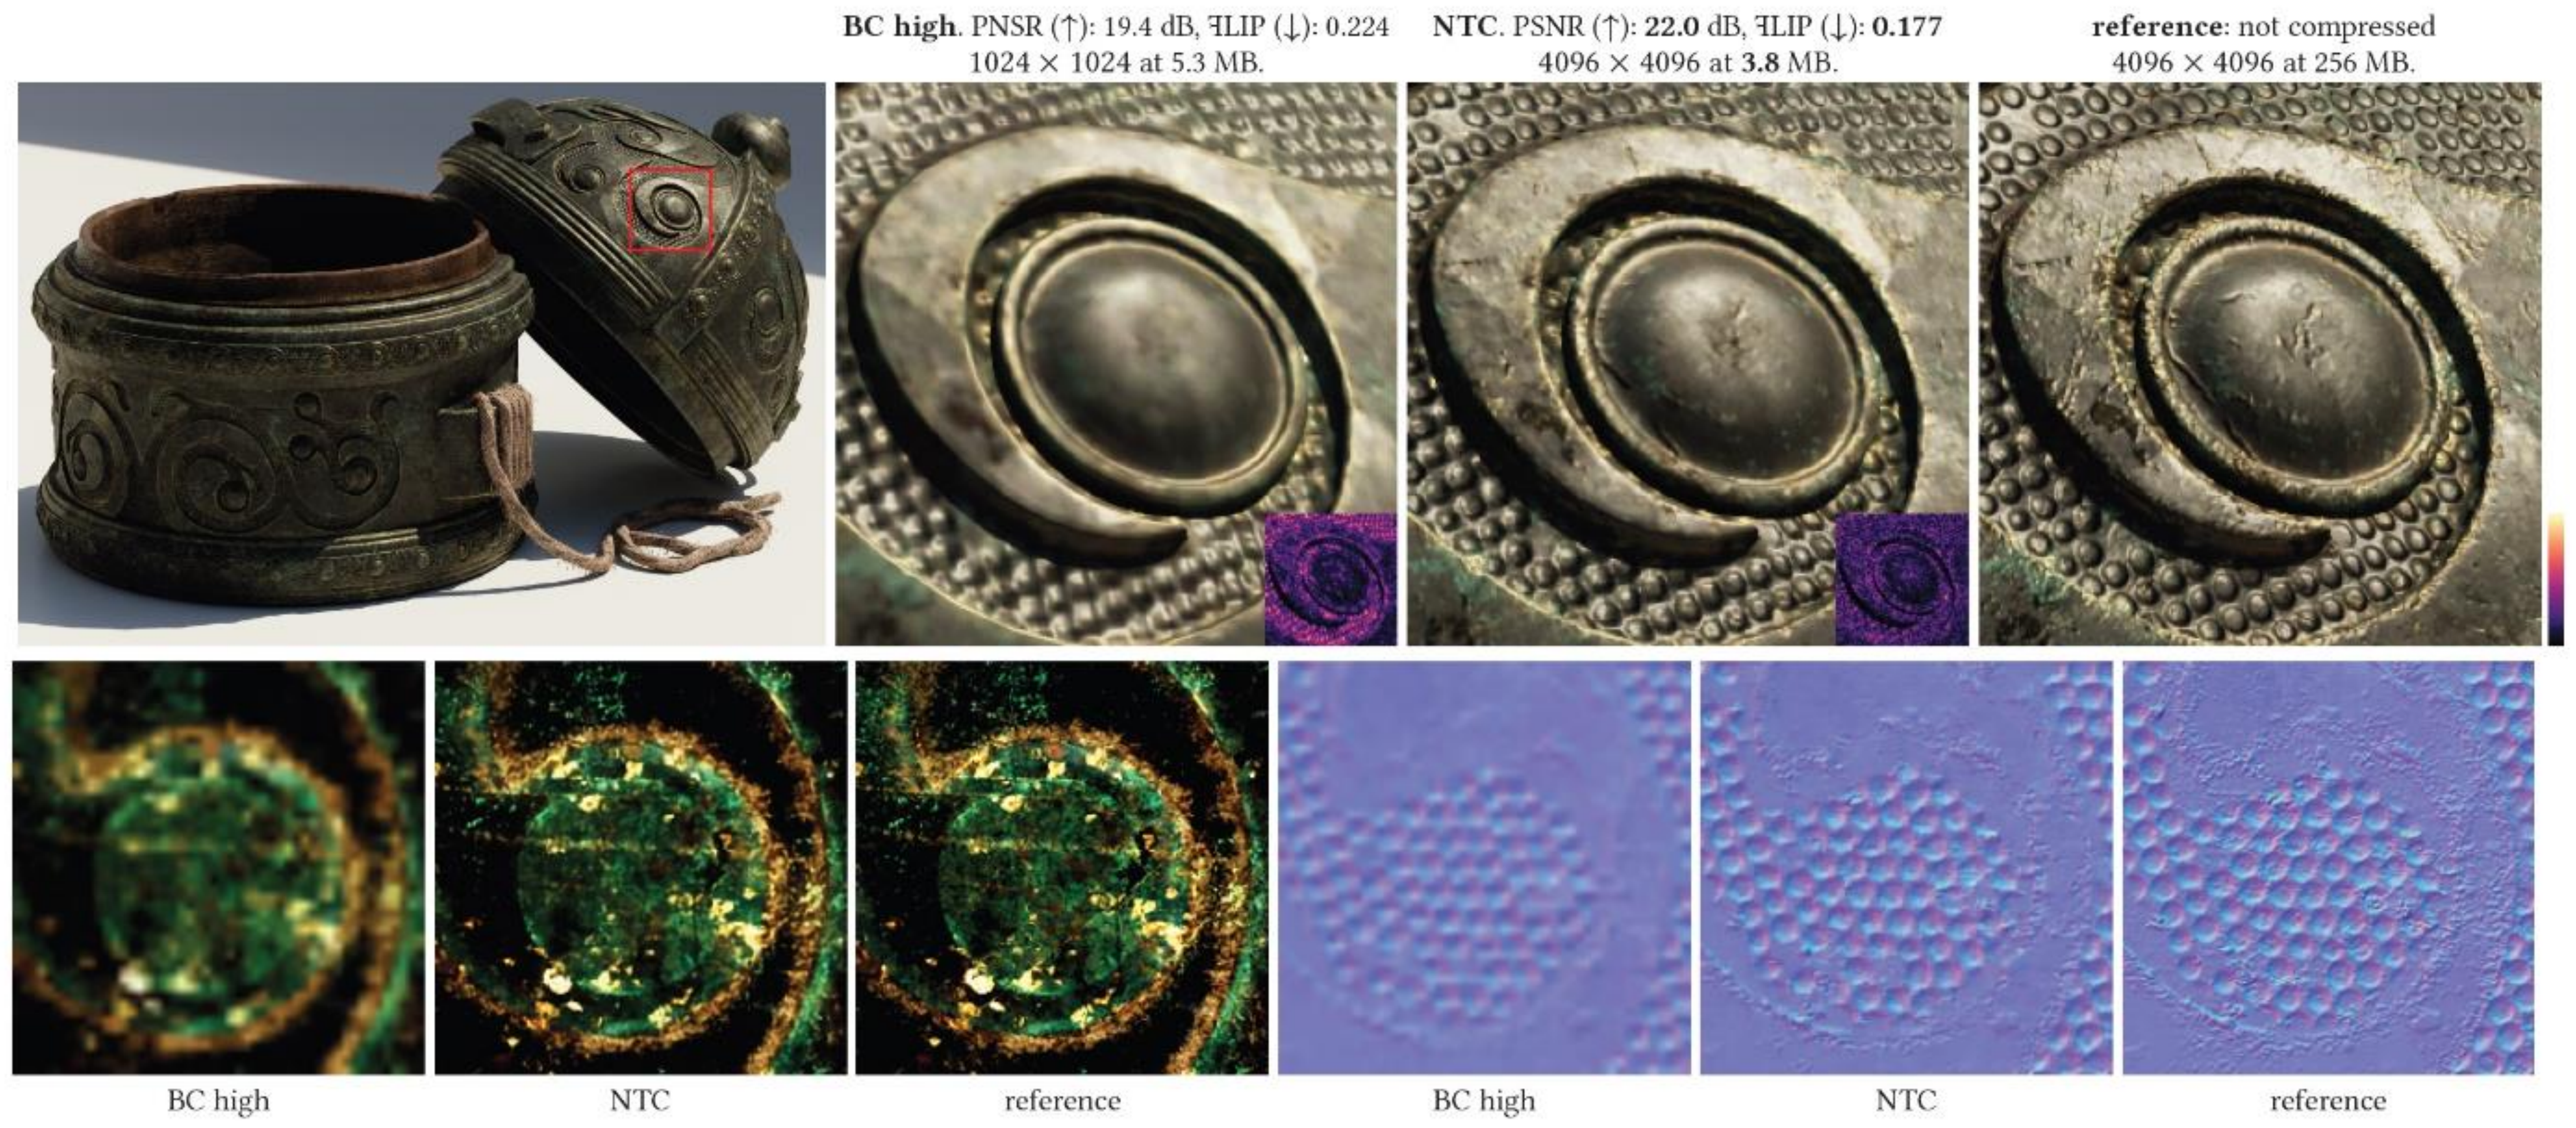

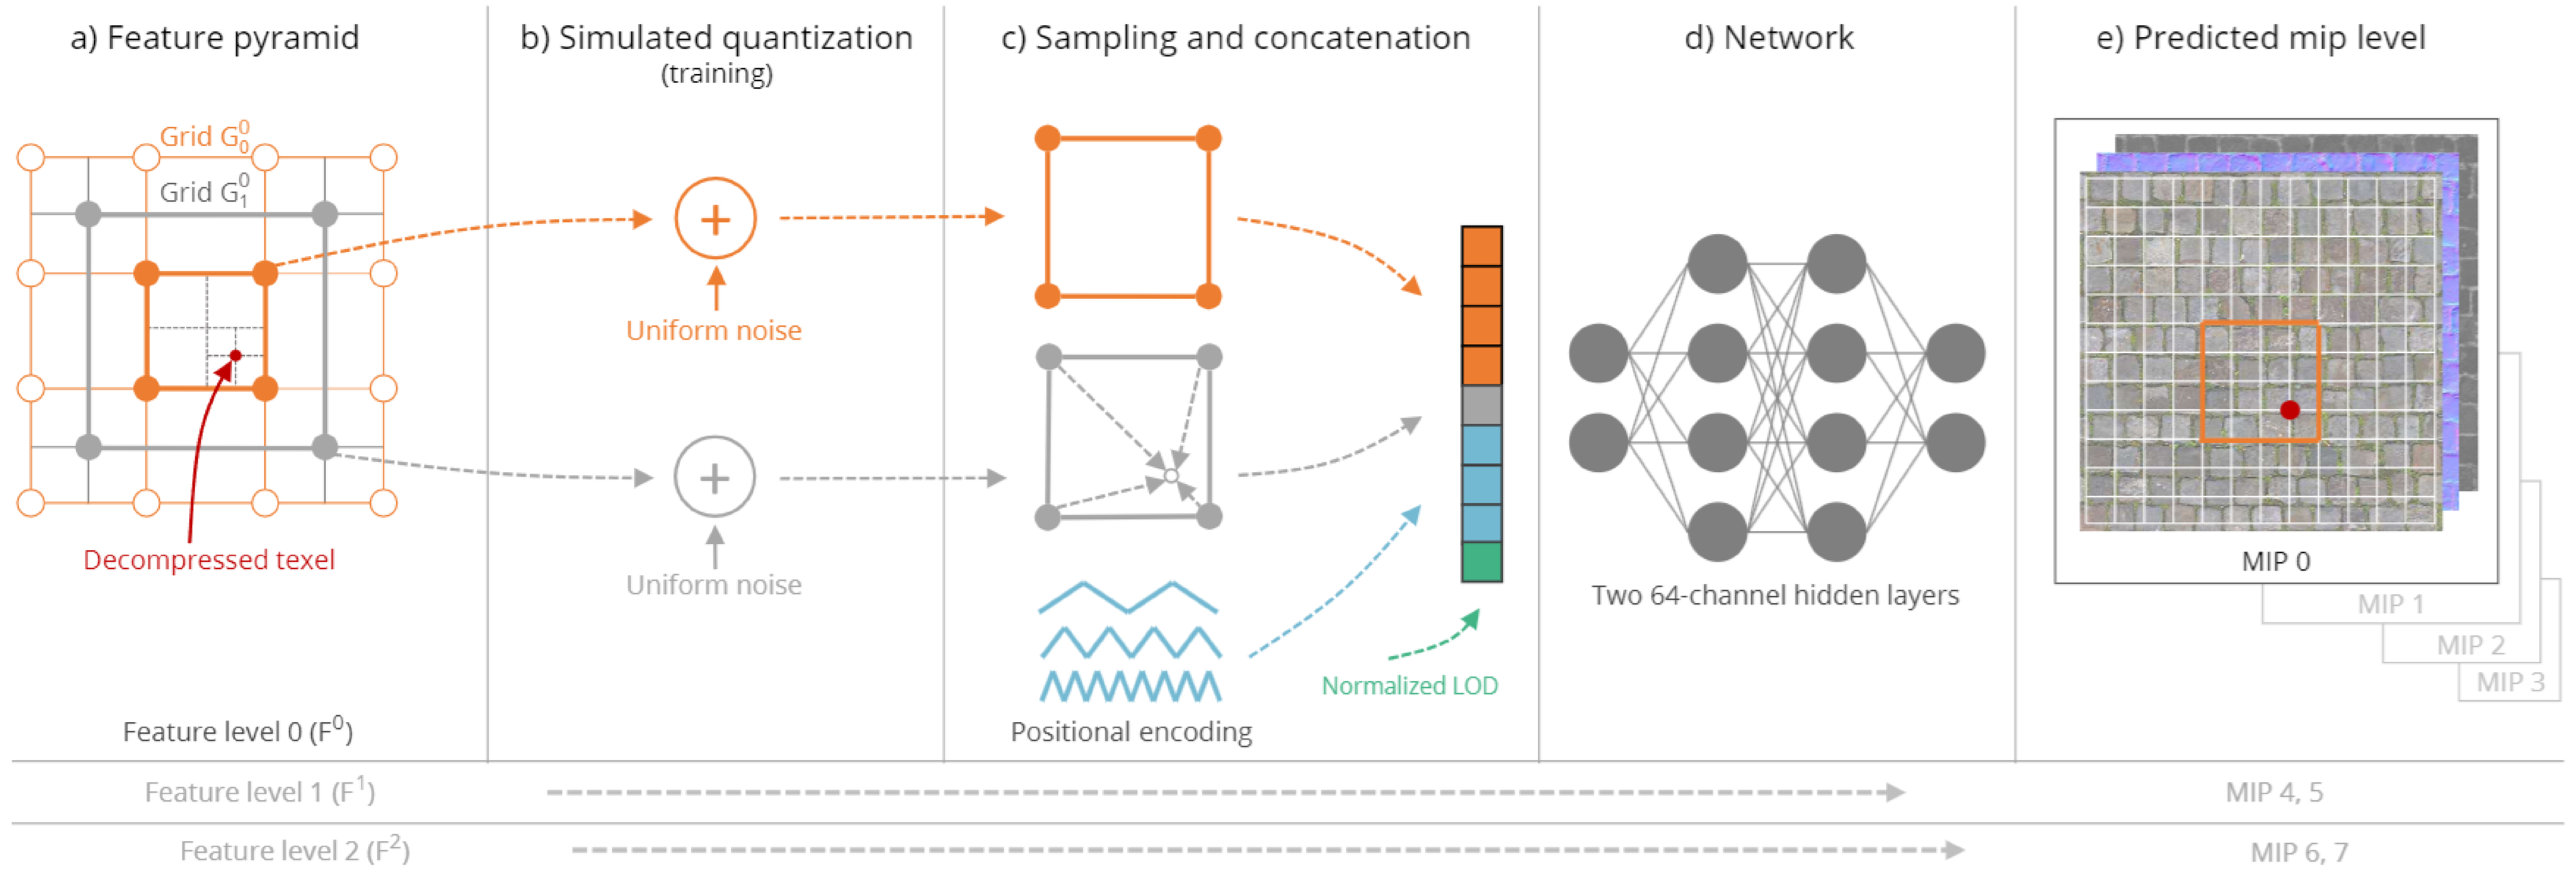

Random-Access Neural Compression of Material Textures, Vaidyanathan et al., Siggraph 2023

Project beginning : Stochastic neural texture filtering (performance)

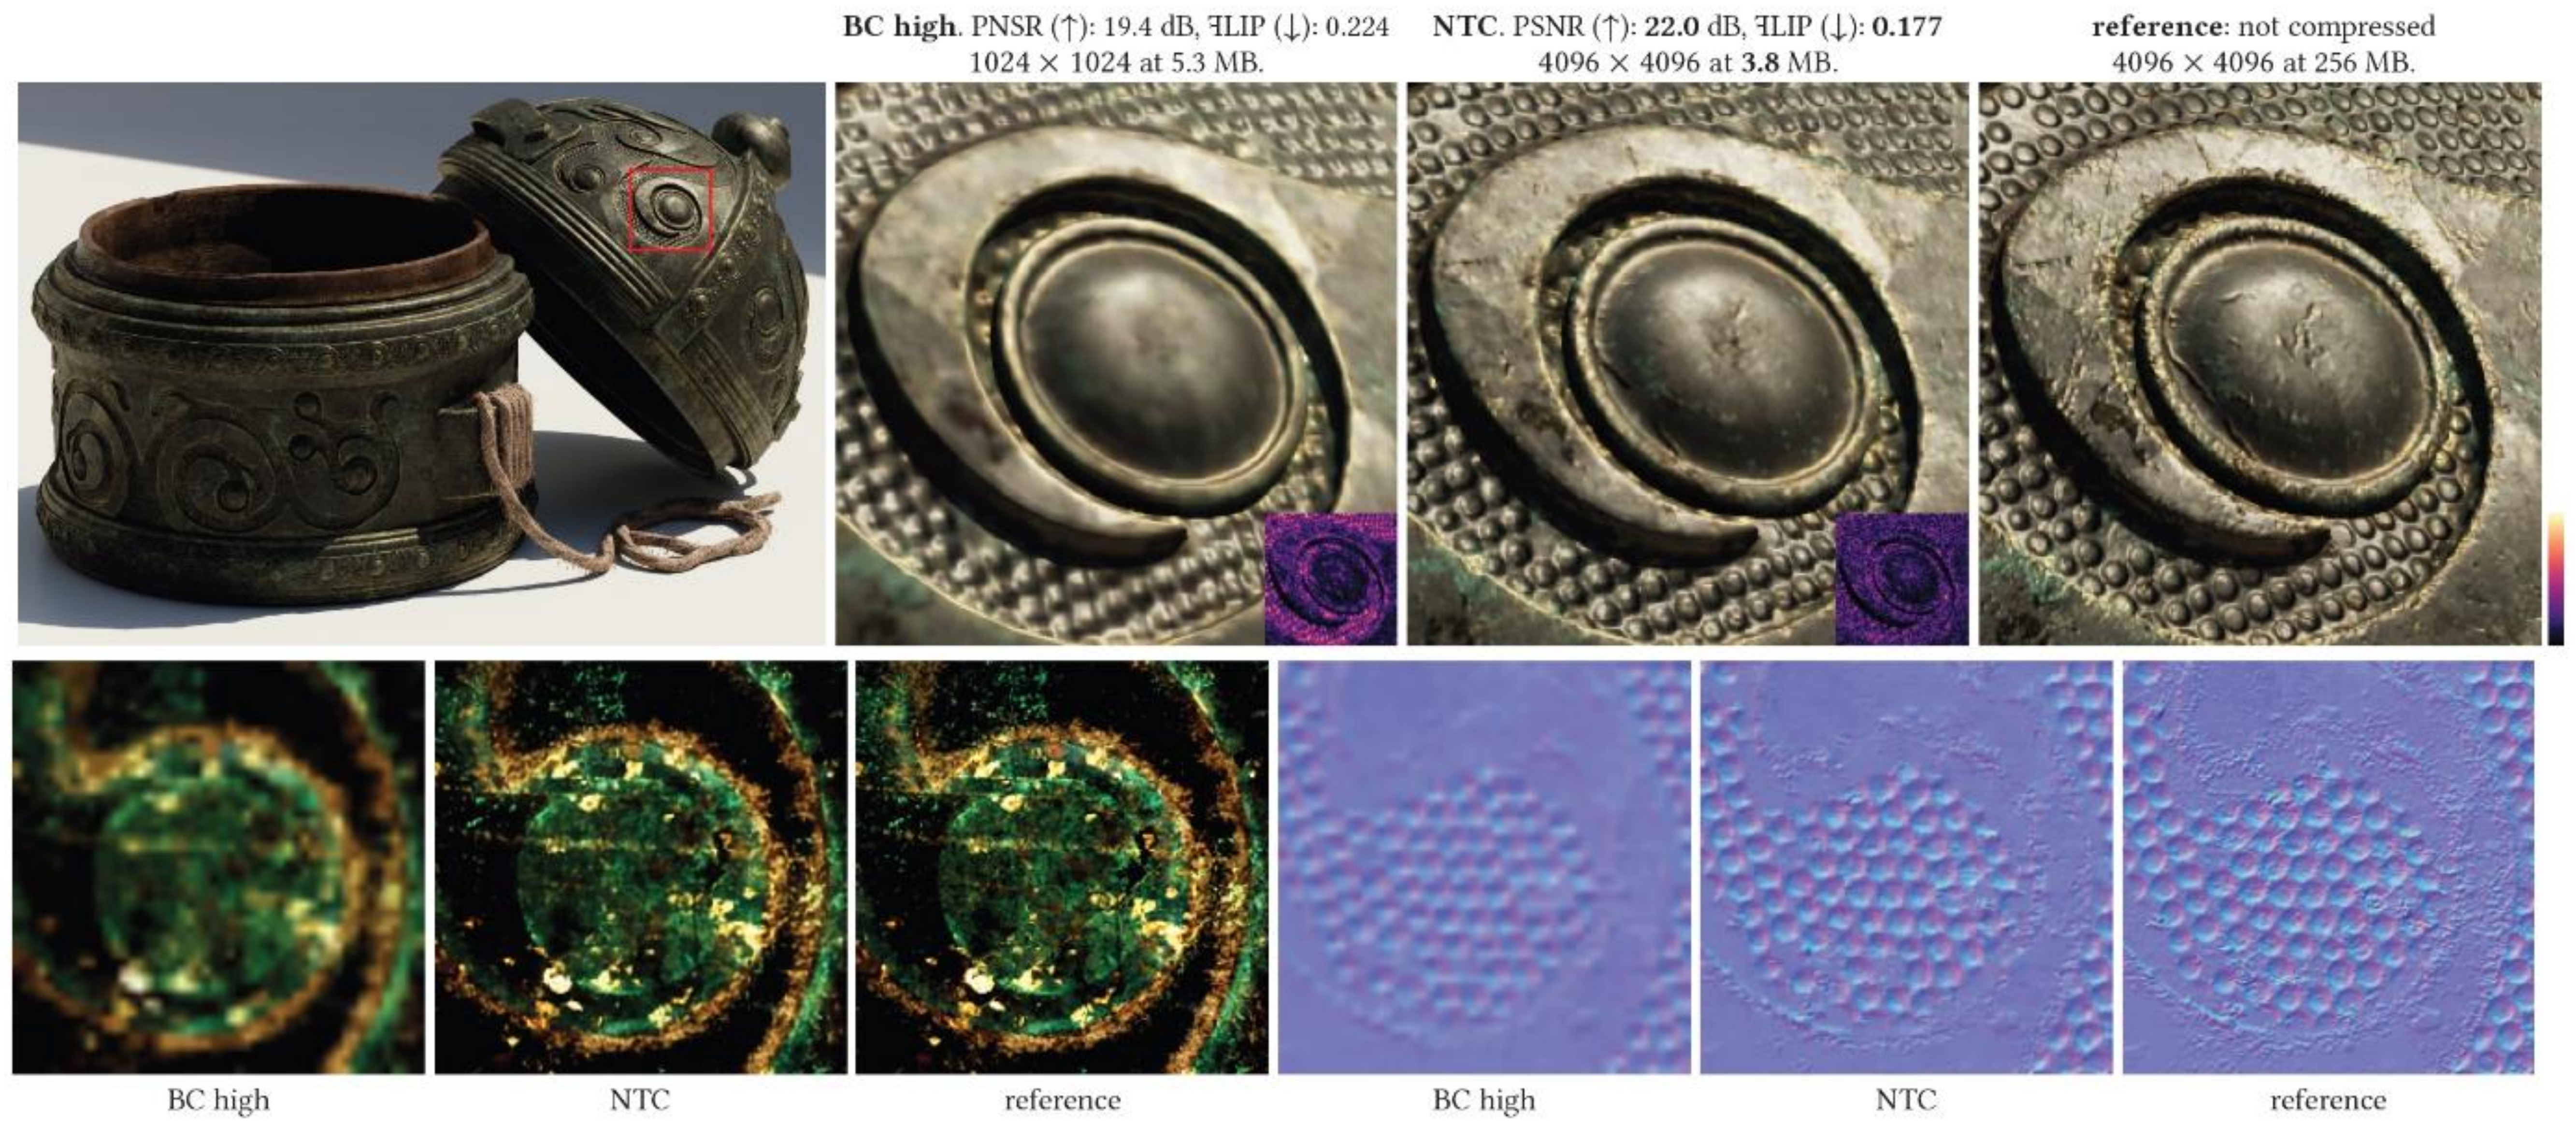

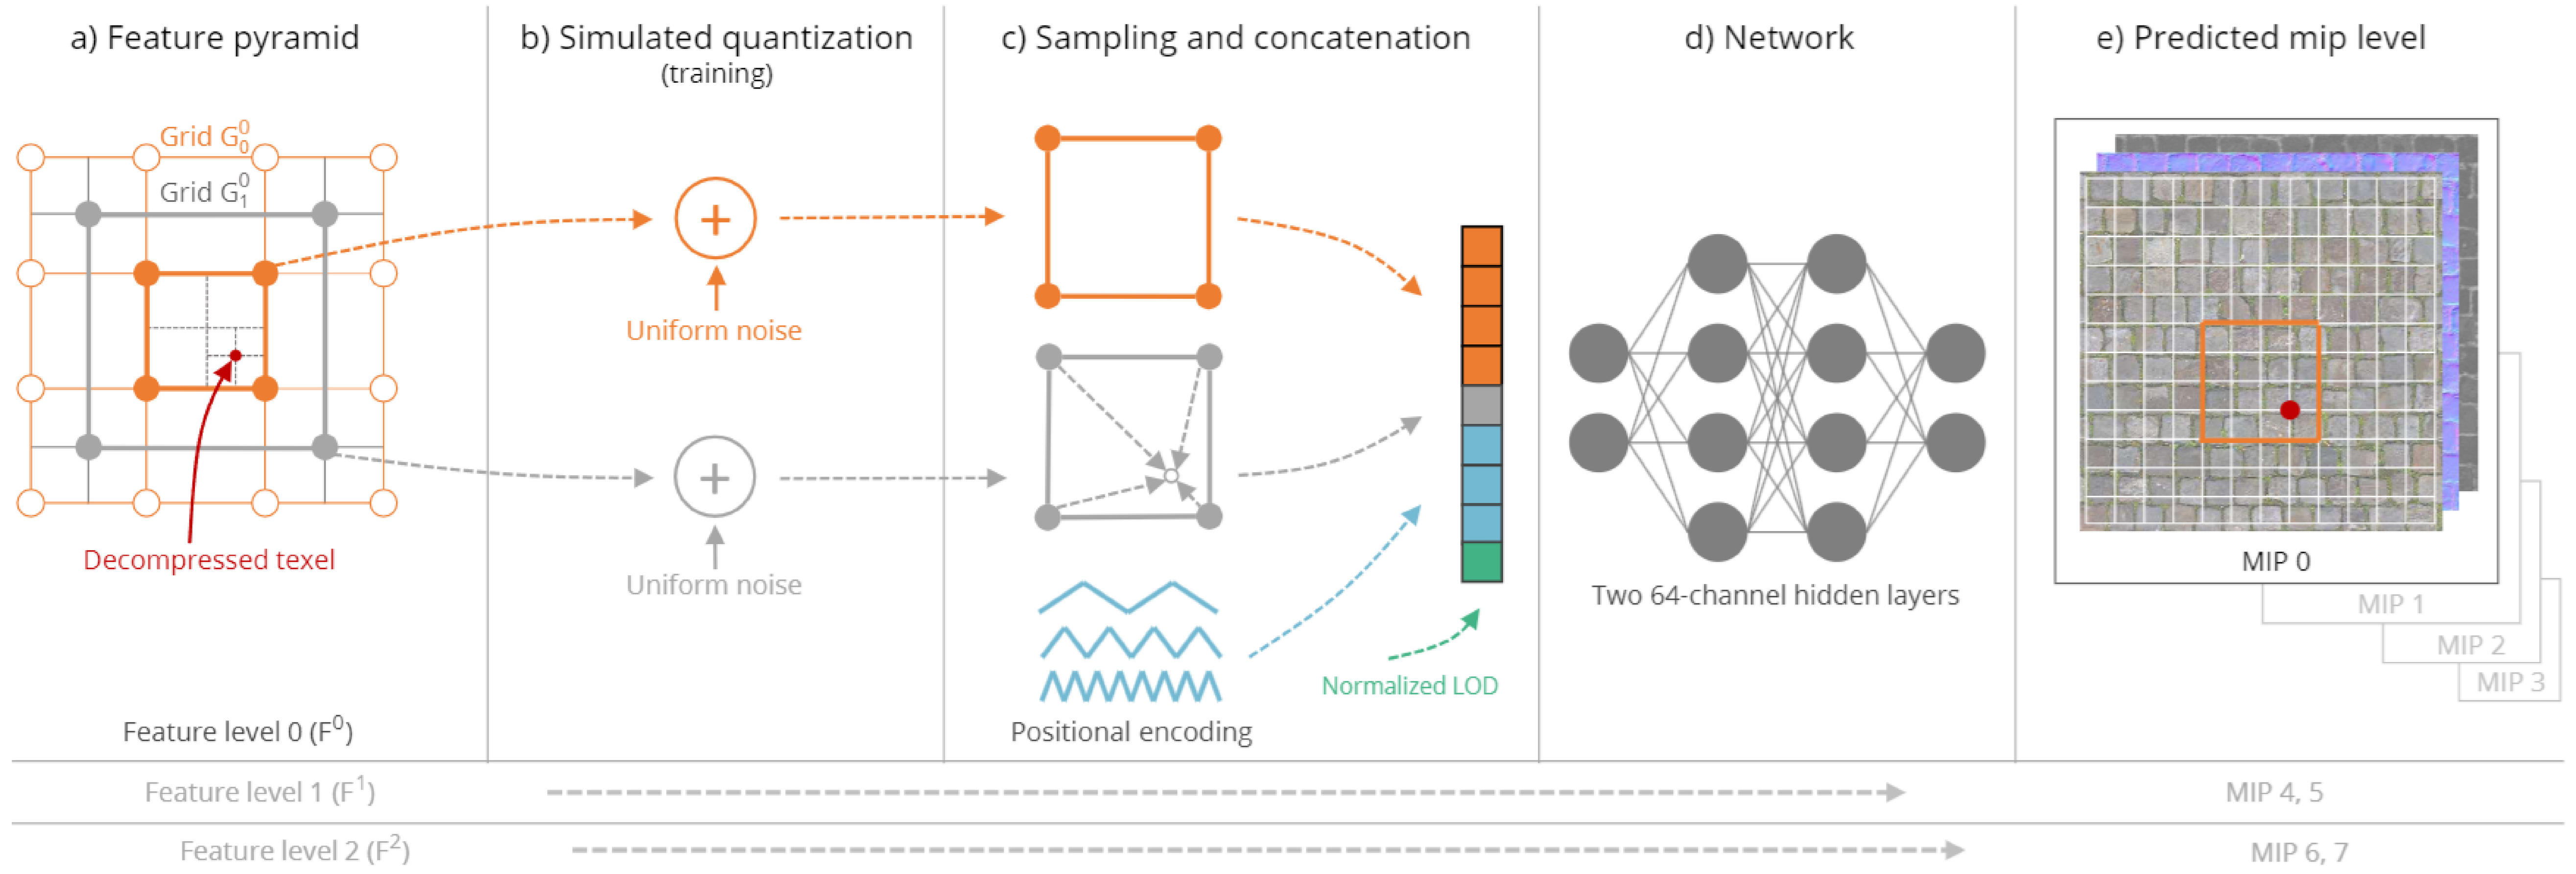

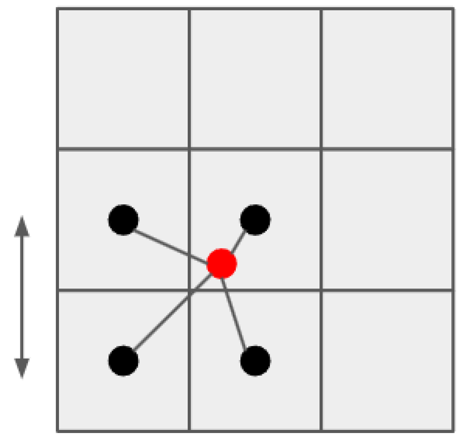

Random-Access Neural Compression of Material Textures, Vaidyanathan et al., Siggraph 2023

Project beginning : Stochastic neural texture filtering (performance)

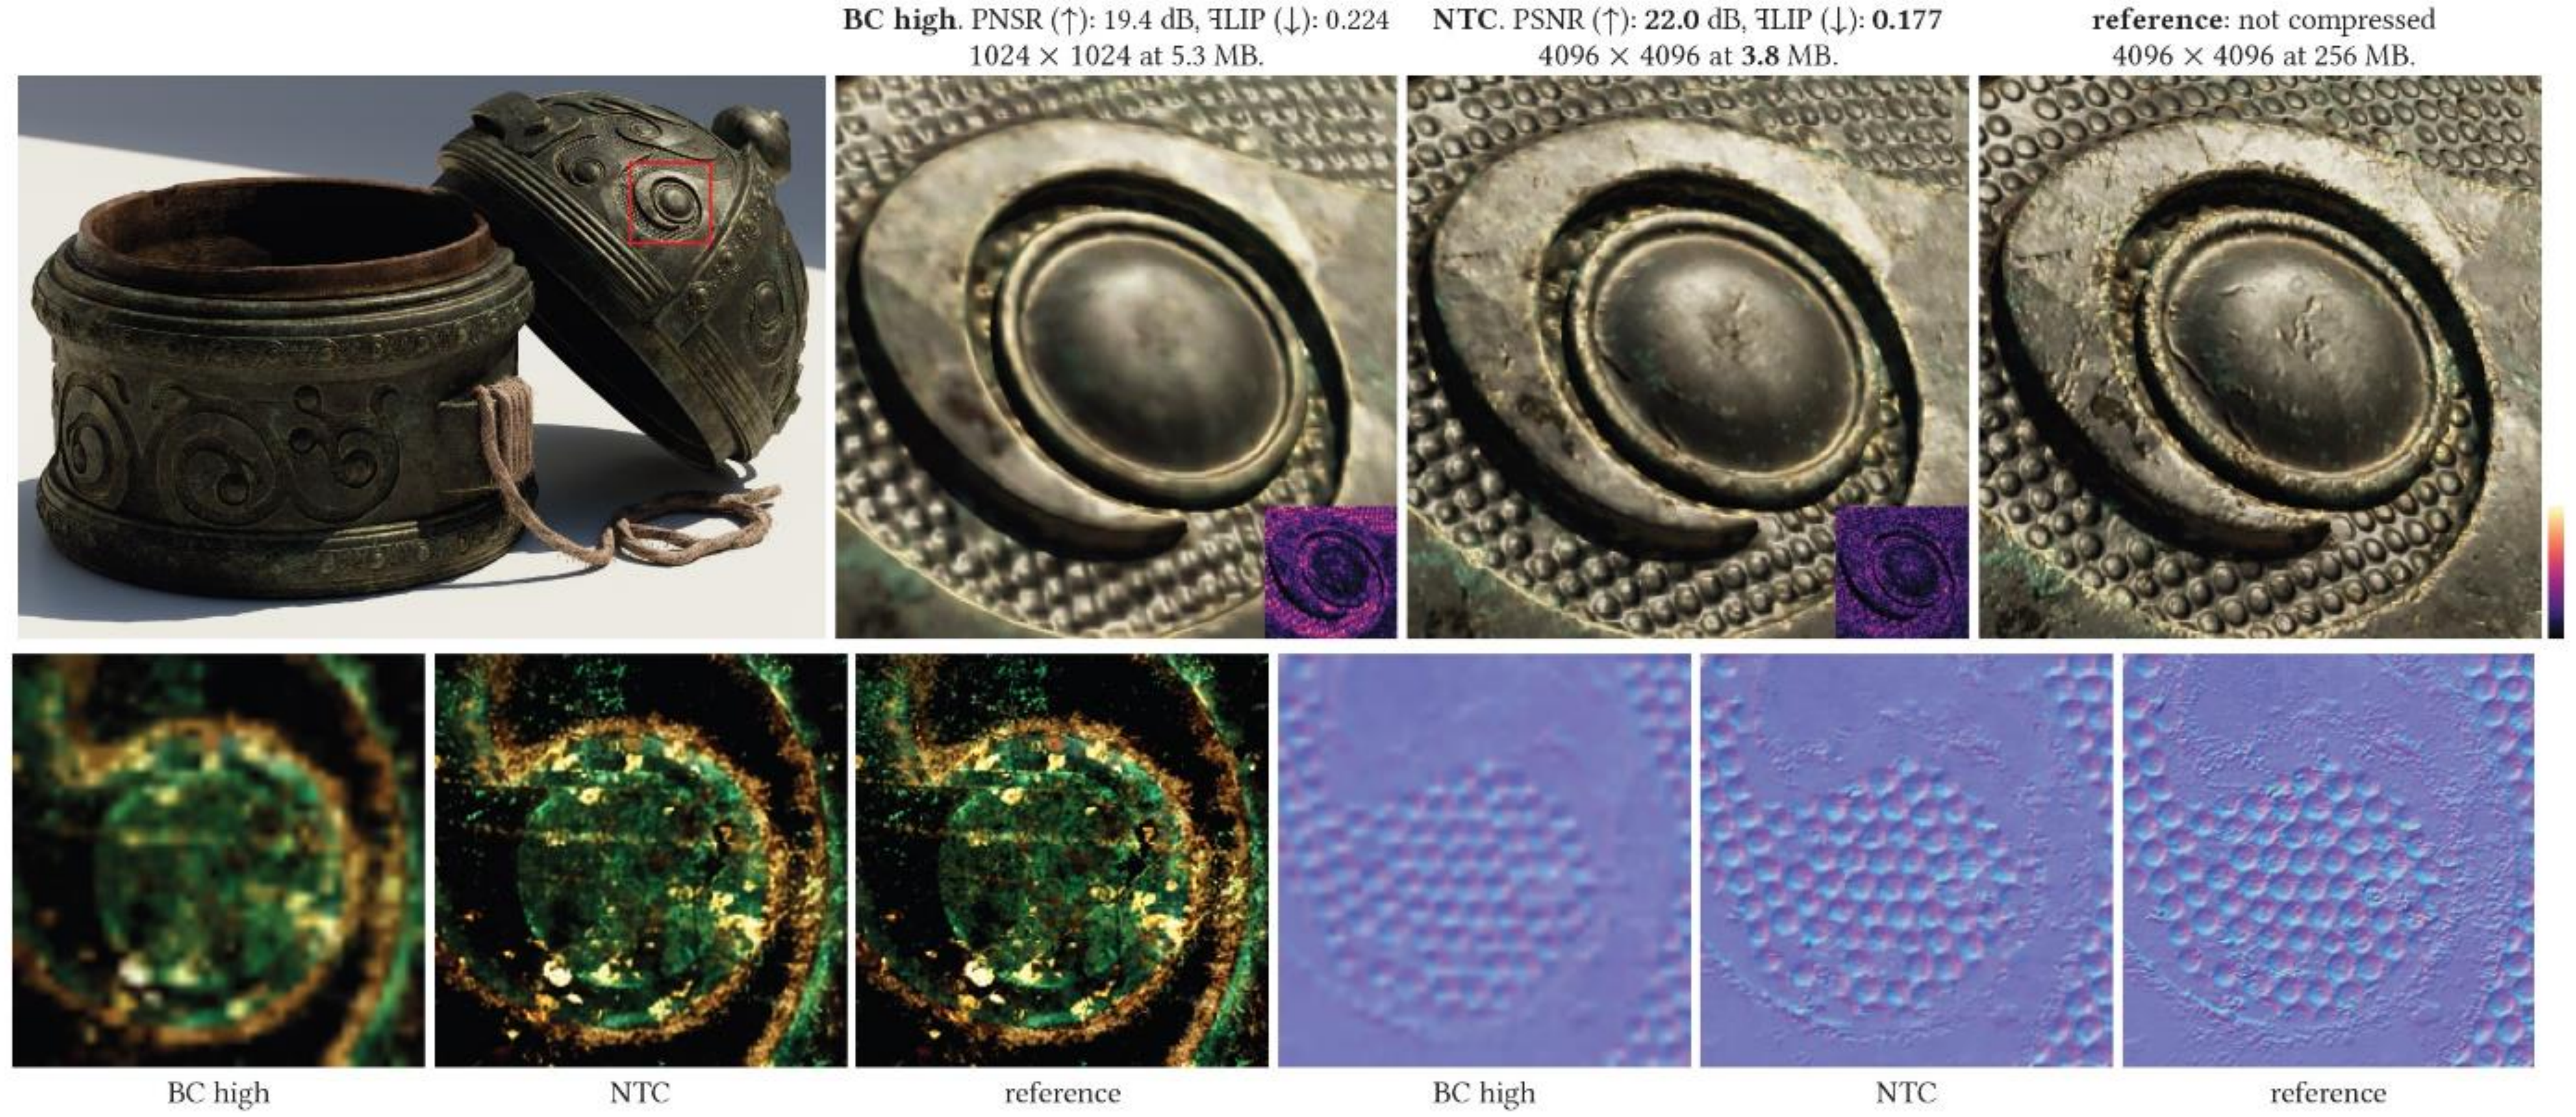

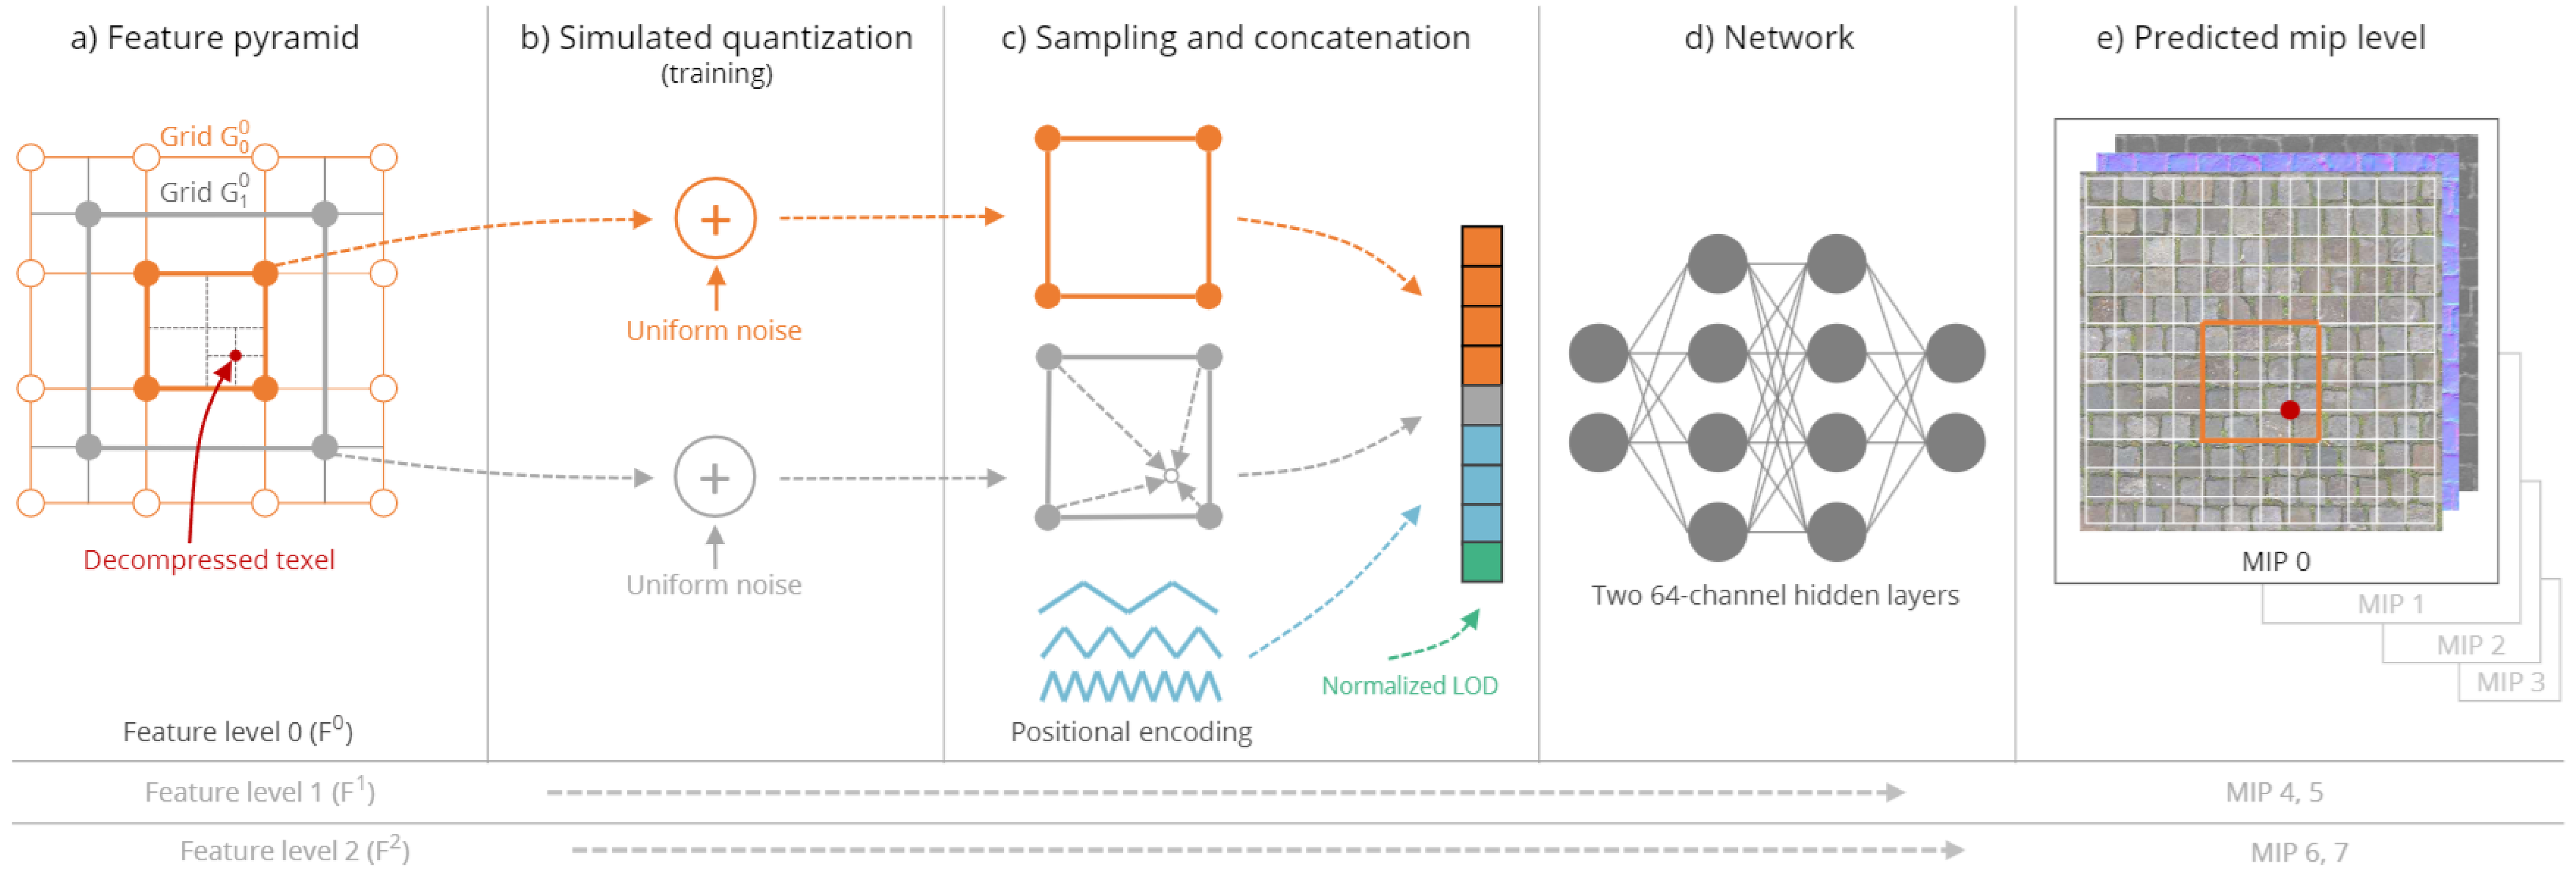

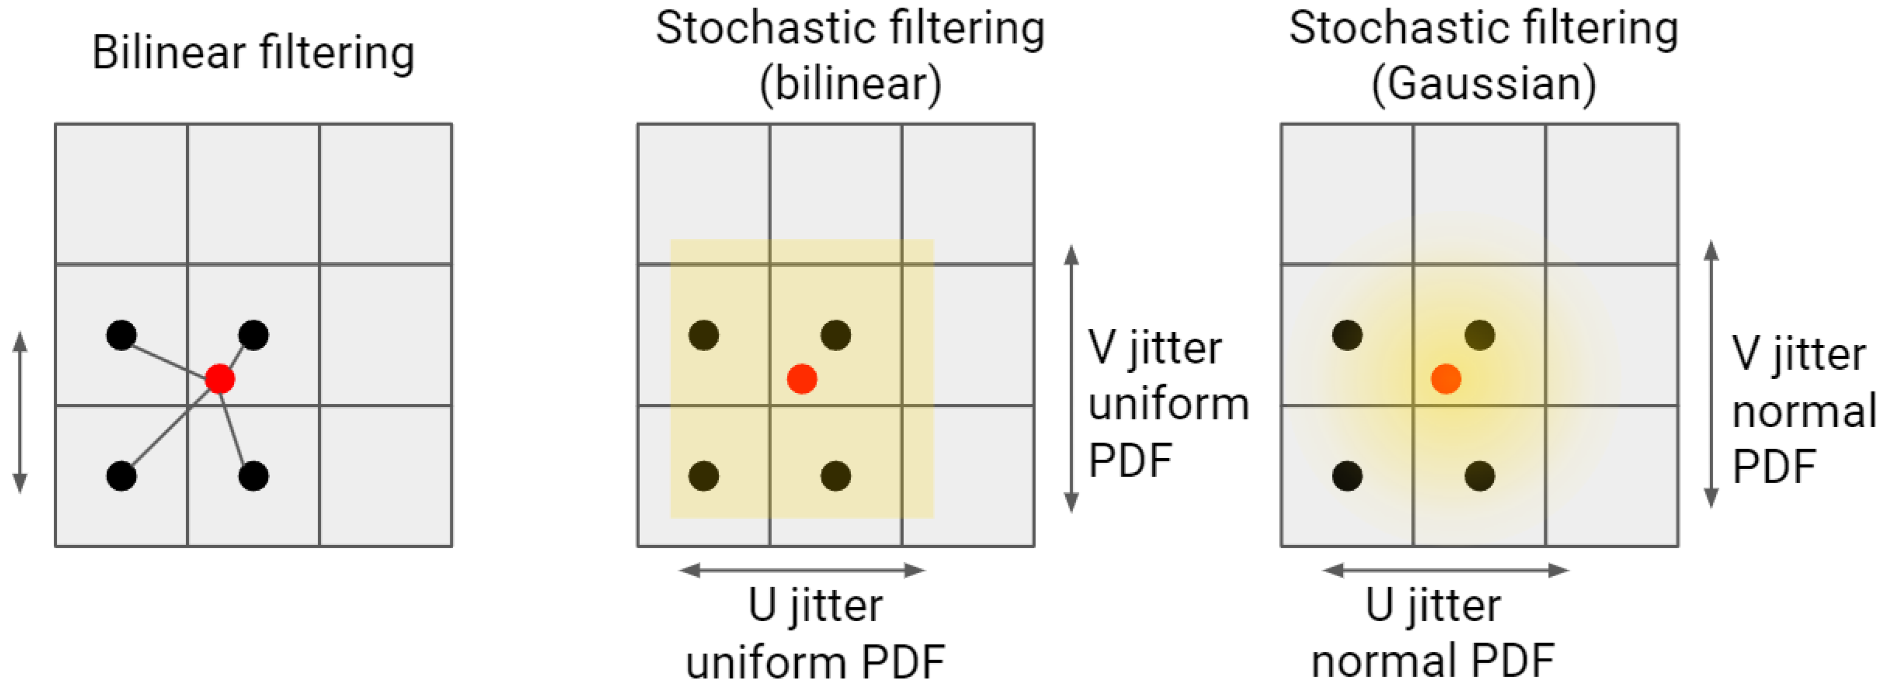

***Random-Access Neural Compression of Material Textures*, Vaidyanathan et al., Siggraph 2023**

Related Work – Long History of Stochastic Filtering

- Percentage closer filtering
- Original UE software rasterizer: texture-space dithered nearest lookups
- Star Trek, 25th Anniversary: dithered bilinear filtering
- Negative LOD biasing (UE + everyone using TAA/DLSS...)
- *OpenImageIO*: stochastic LOD selection (via Max Liani)
- Dreamworks *MoonRay*: nearest sampling for minification, bilerp for magnification
- *Interactive Path Tracing and Reconstruction of Sparse Volumes*, Hofmann, Hasselgren, Clarberg, and Munkberg, i3d 2021: stochastic trilinear
- *Random-Access Neural Compression of Material Textures*, Vaidyanathan, Salvi, Wronski, Akenine-Möller, Ebelin, Lefohn, SIGGRAPH 2023: stochastic trilinear

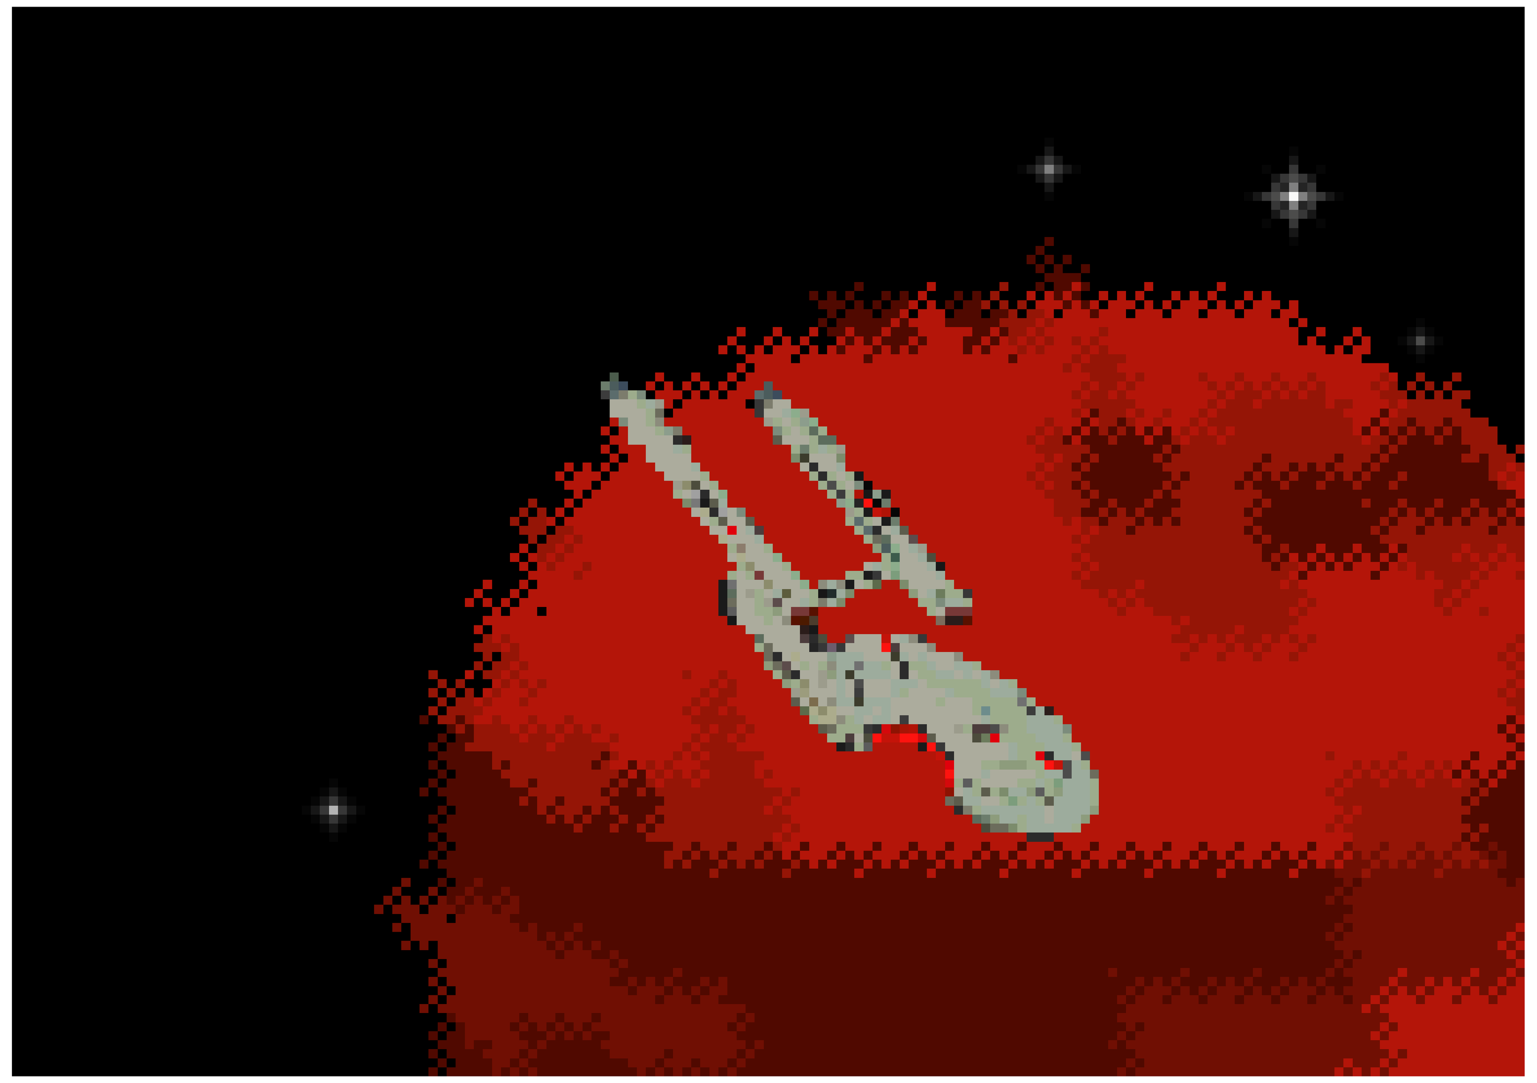

Related Work – Long History of Stochastic Filtering

- Percentage closer filtering
- Original UE software rasterizer: texture-space dithered nearest lookups
- Star Trek, 25th Anniversary: dithered bilinear filtering
- Negative
- *OpenIma*
- Dreamw
  magnific
- *Interacti*
  Hofmanr
  trilinear
- *Random*
  Vaidyanathan, Salvi, Wronski, Akenine-Möller, Ebelin, Lefohn, SIGGRAPH 2023: stochastic trilinear

Advances in temporal reconstruction make stochastic techniques very attractive

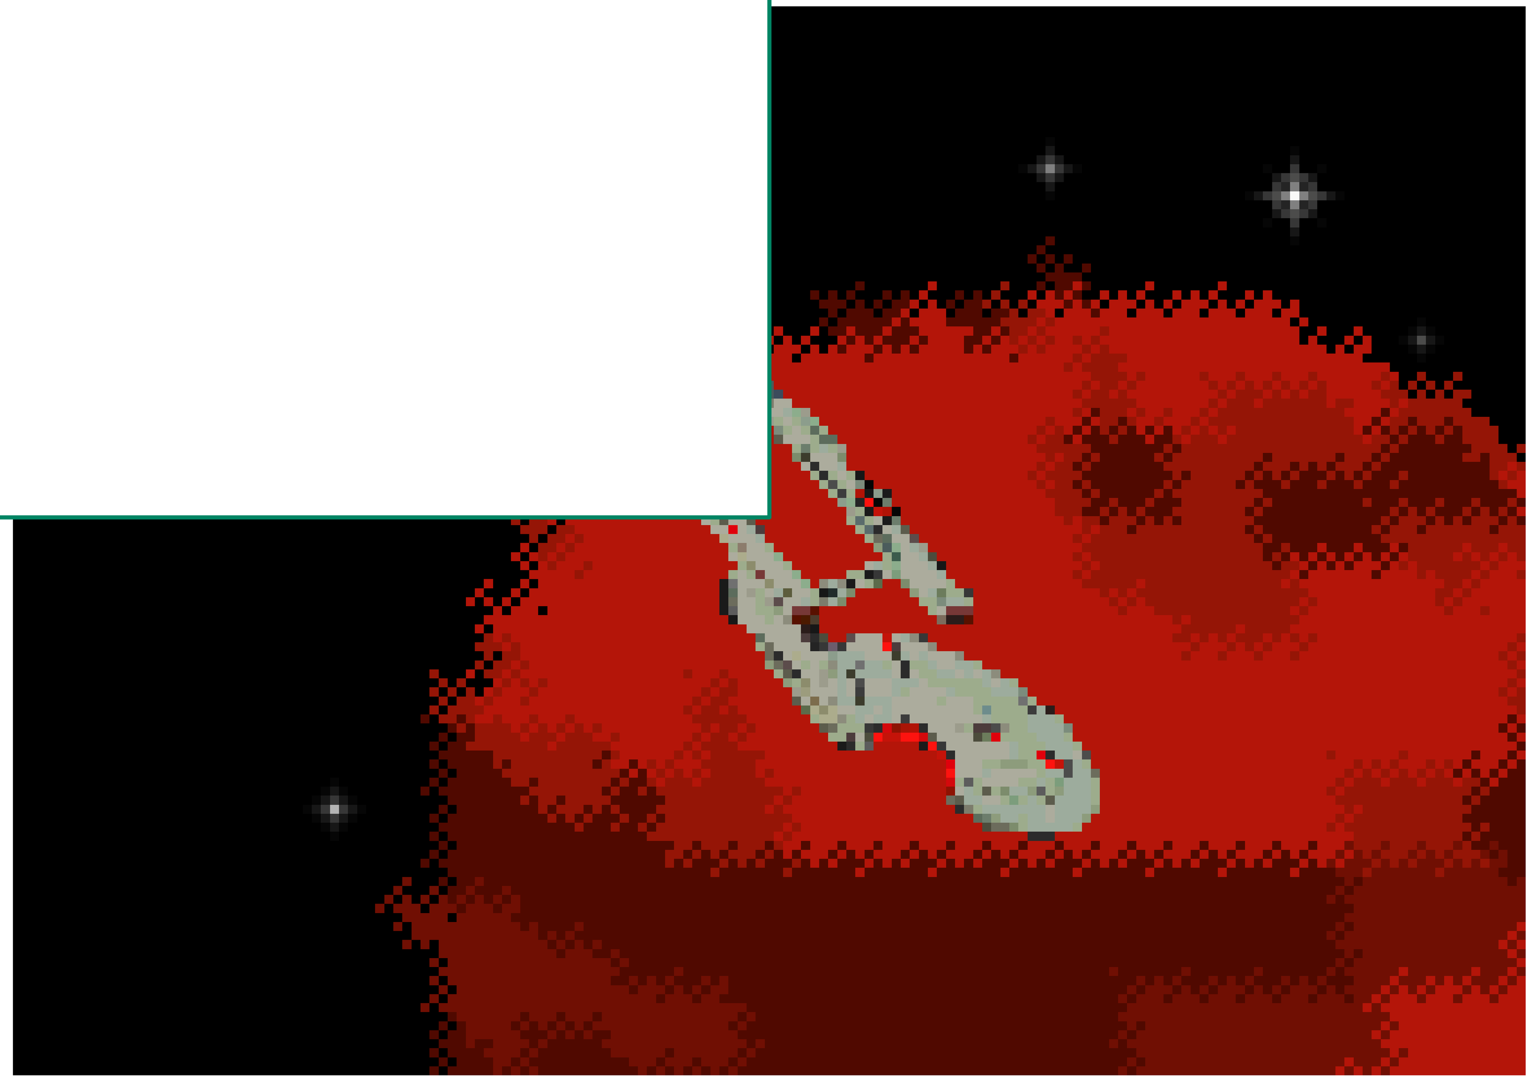

Related Work – Long History of Stochastic Filtering

- Percentage closer filtering
- Original UE software rasterizer: texture-space dithered nearest lookups
- Star Trek 25th Anniversary: dithered bilinear filtering
- Negative
- *OpenIma*
- Dreamw
  magnific
- *Interacti*
  Hofmann
  trilinear
- *Random*
  Vaidyanathan, Salvi, Wronski, Akenine-Möller, Ebelin, Lefohn, SIGGRAPH 2023: stochastic trilinear

Advances in temporal reconstruction make stochastic techniques very attractive

We will generalize those, formalize, and propose two families of techniques beyond simple filters

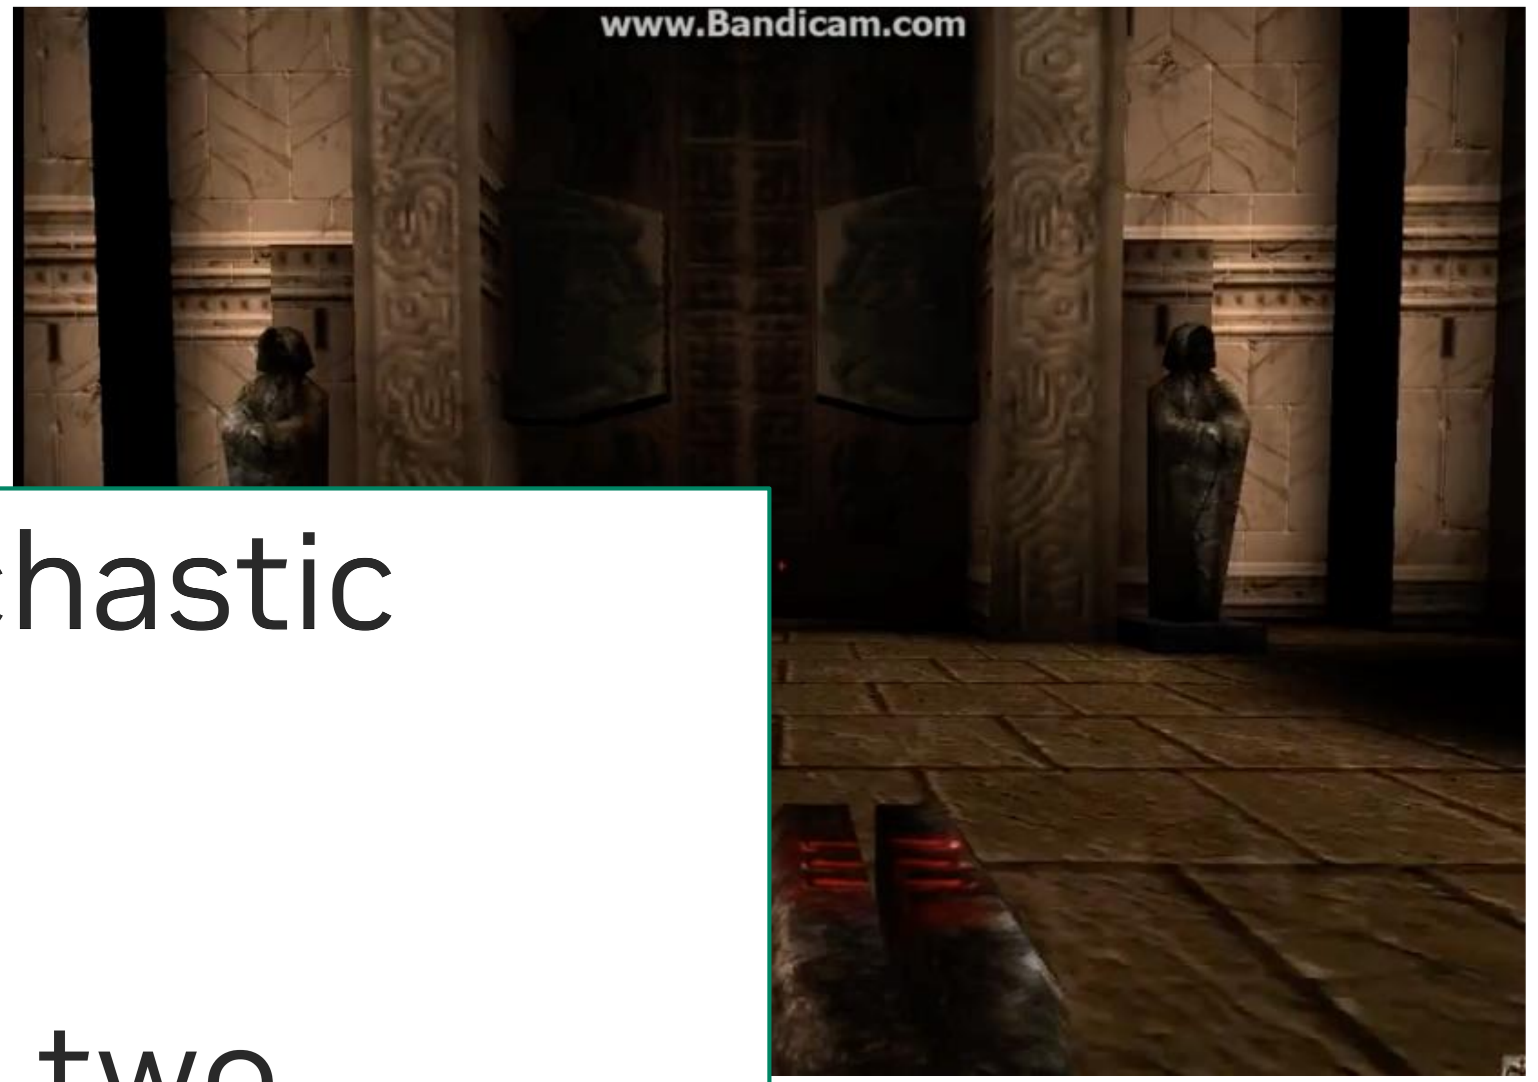

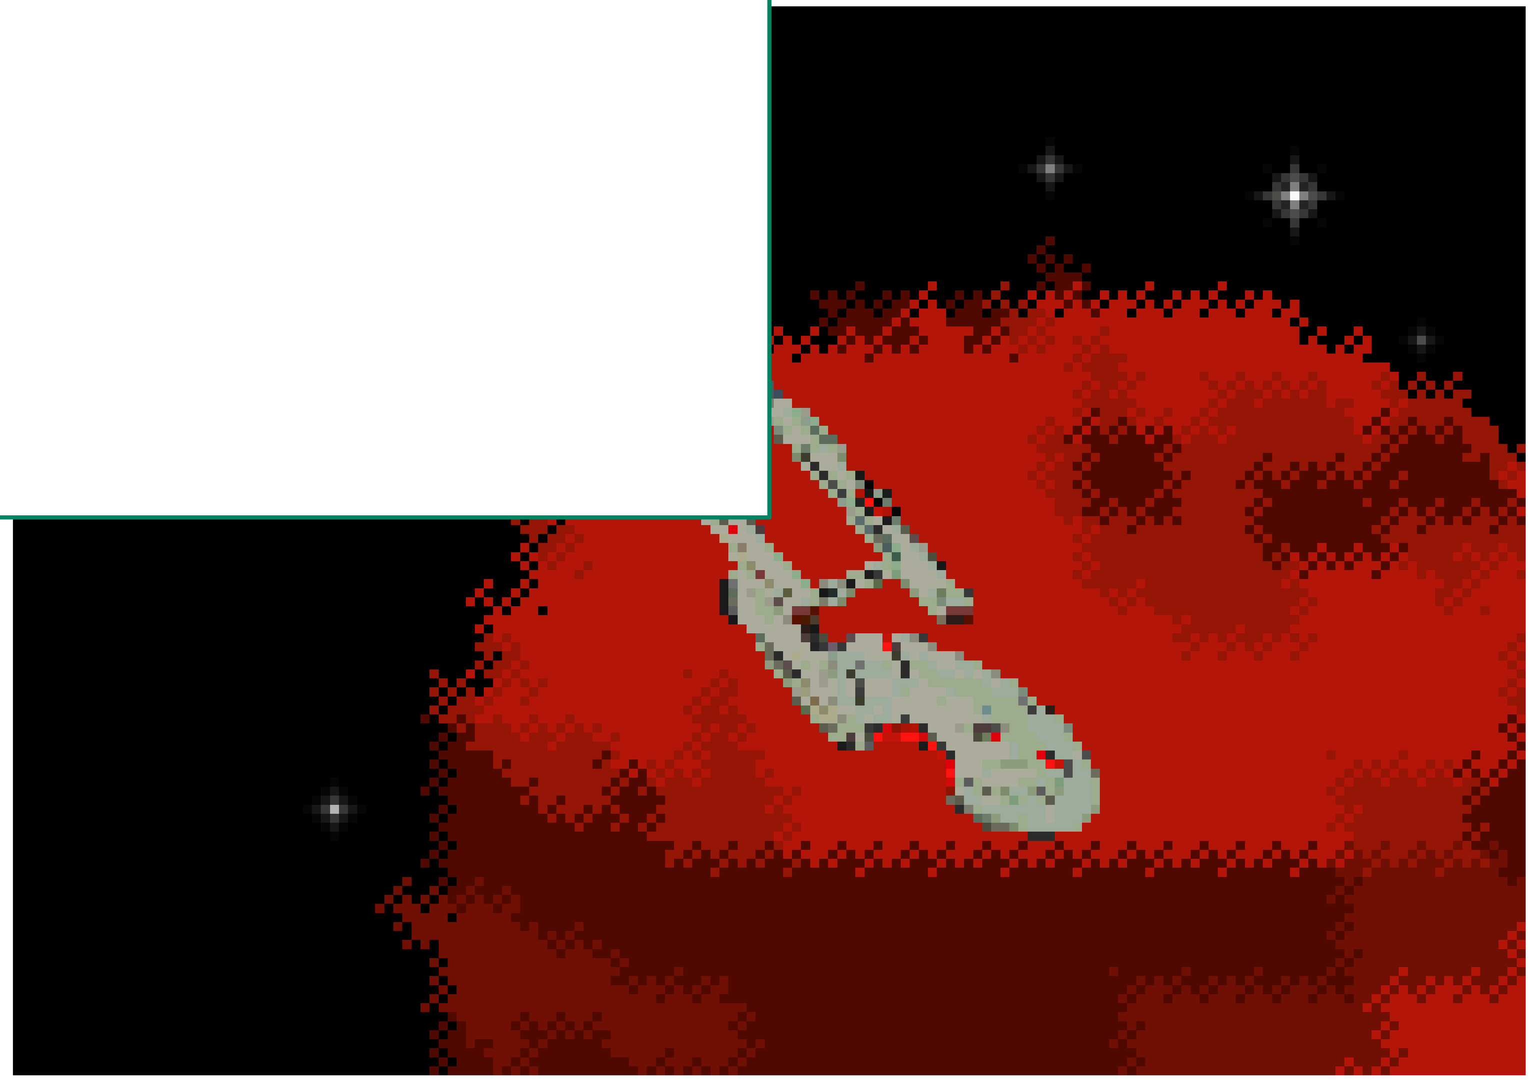

Related Work – Long History of Stochastic Filtering

- Percentage closer filtering
- Original UE software rasterizer: texture-space dithered nearest lookups
- Star Trek, 25th Anniversary: dithered bilinear filtering
- Negative
- *OpenIma*
- Dreamw magnific
- *Interacti* Hofmanr trilinear
- *Random* Vaidyanathan, Salvi, Wronski, Akenine-Möller, Ebelin, Lefohn, SIGGRAPH 2023: stochastic trilinear

Advances in temporal reconstruction make stochastic techniques very attractive

We will generalize those, formalize, and propose two families of techniques beyond simple filters

But we have a more important problem to solve first…

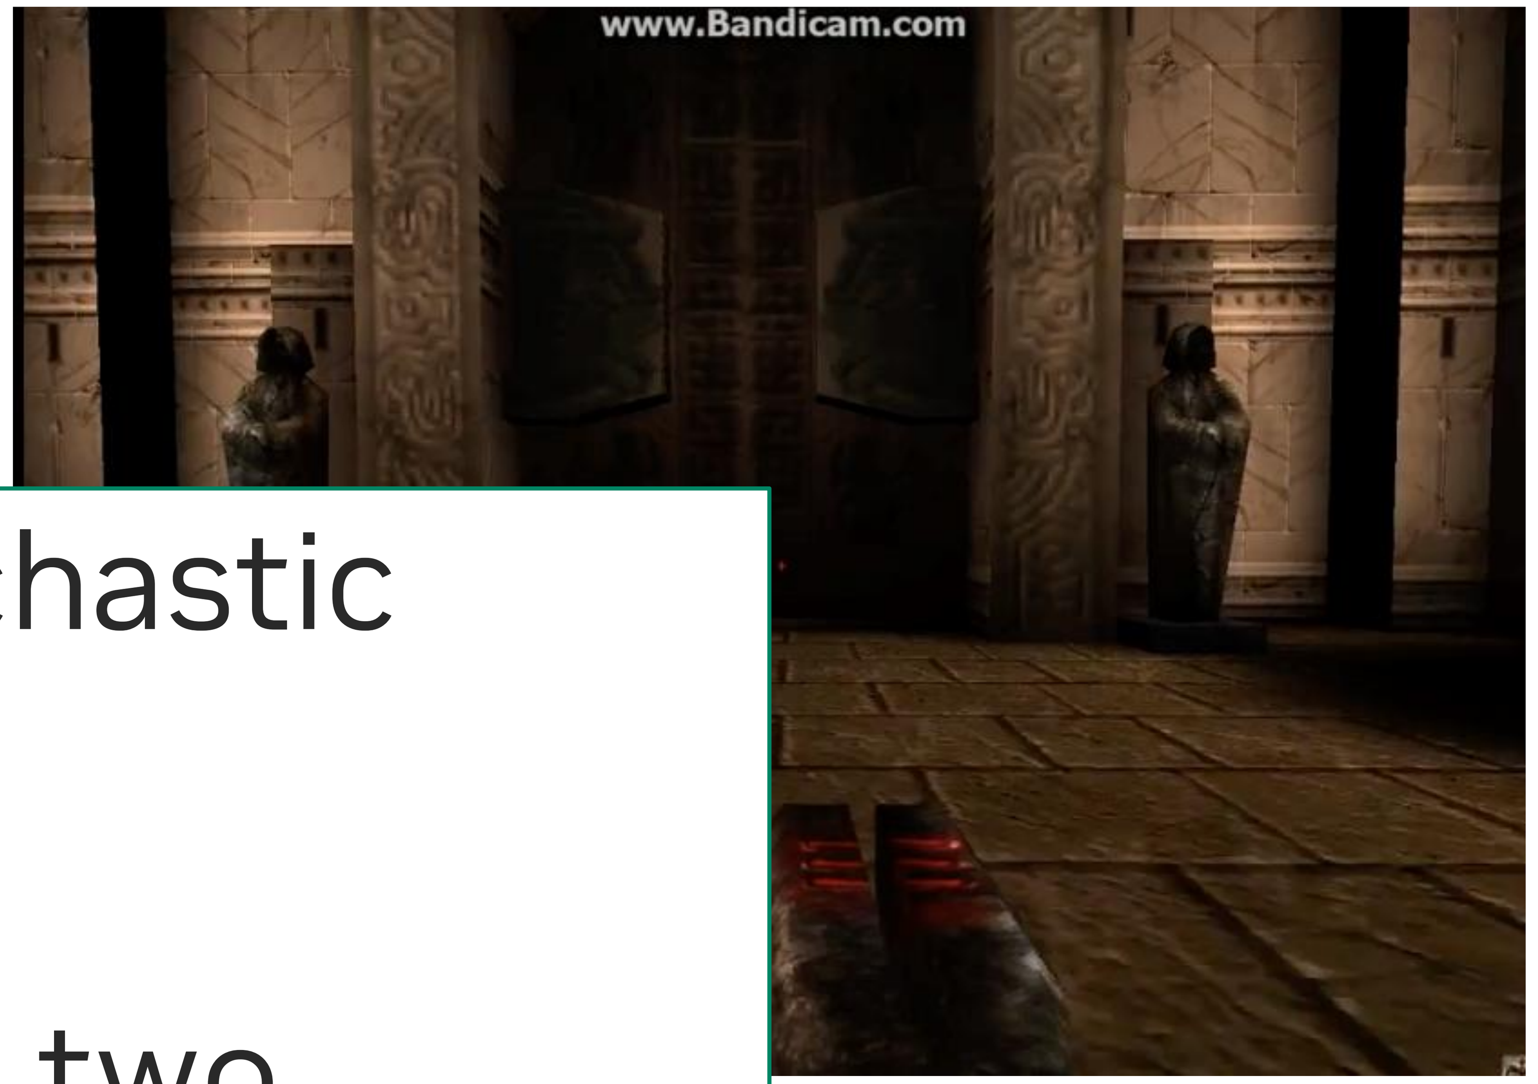

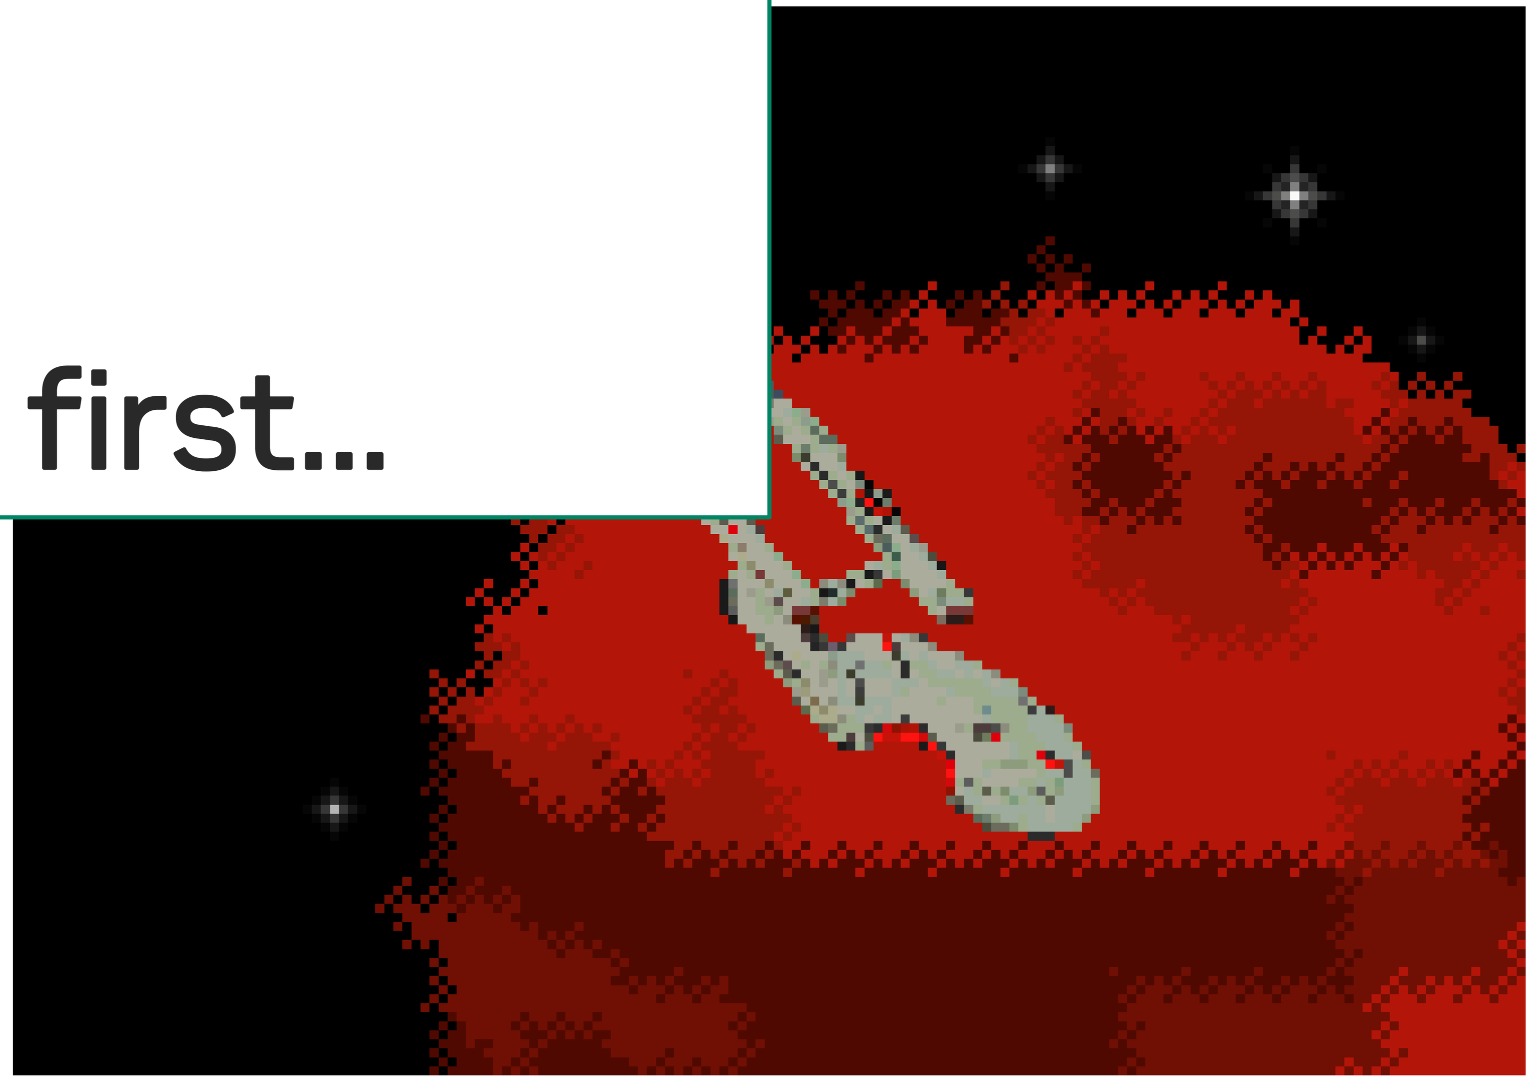

NVIDIA

Stochastic texture filtering: do we have a problem?

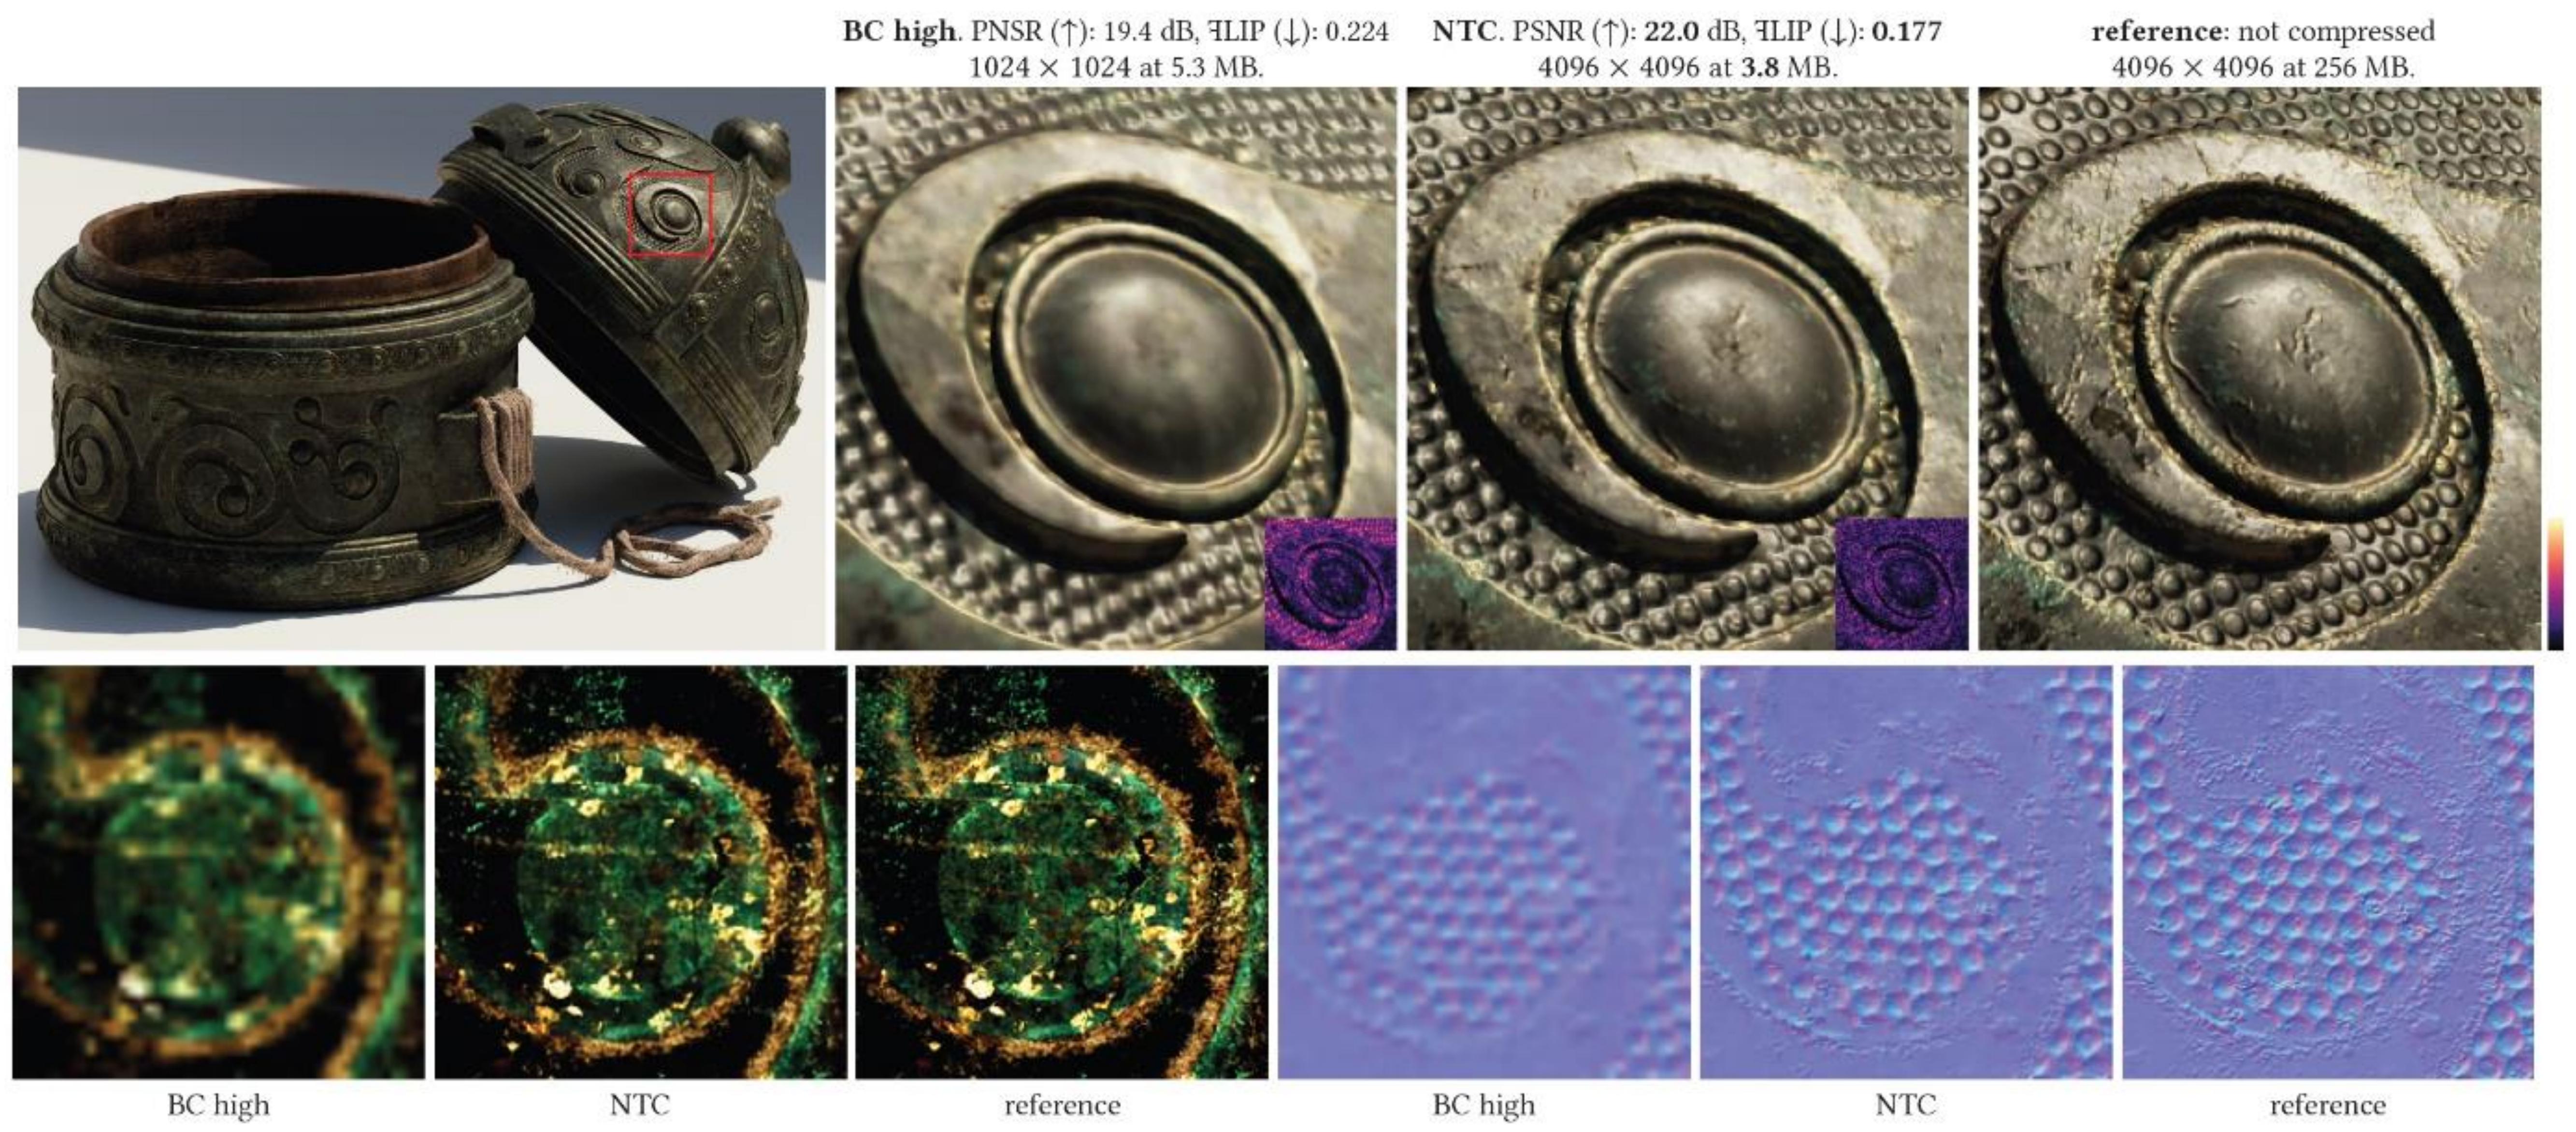

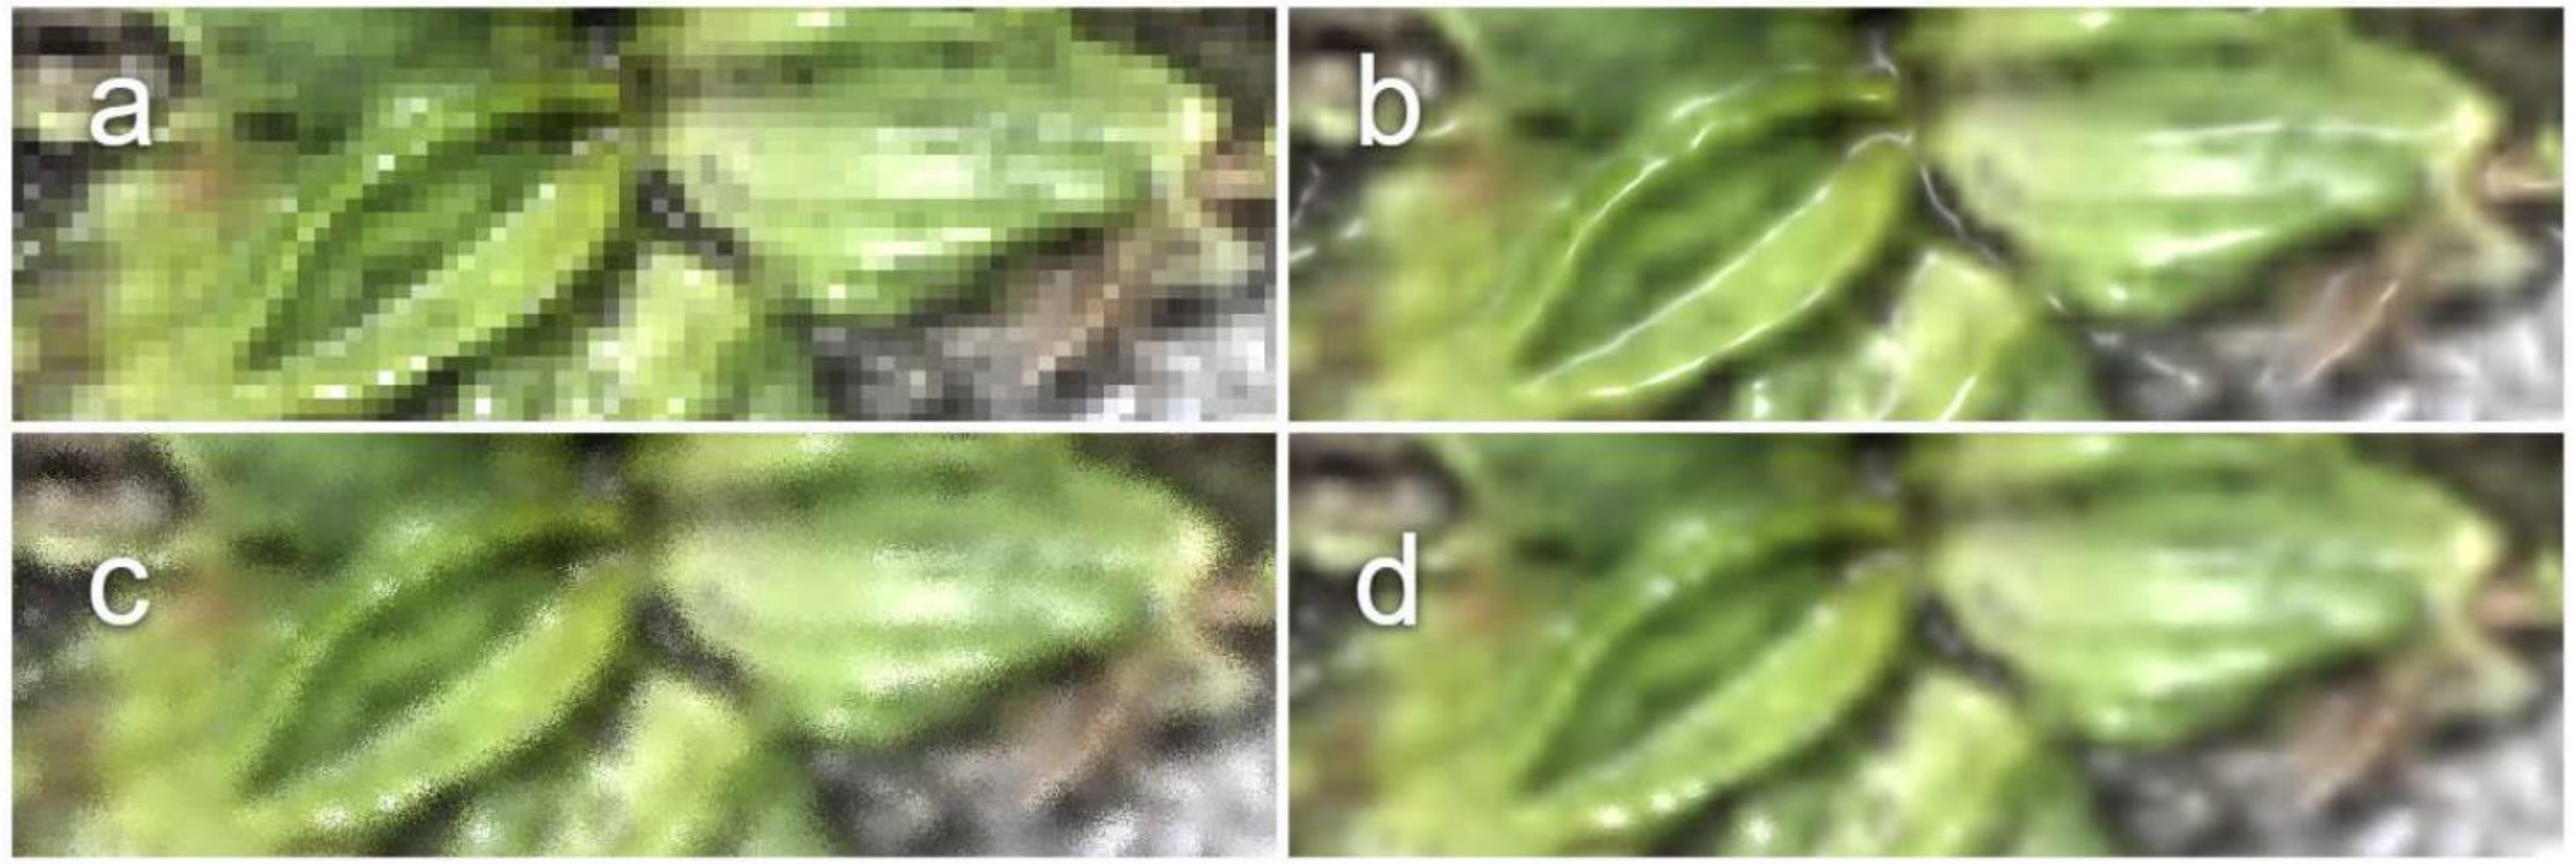

Filtering this way can look significantly different...

Random-Access Neural Compression of Material Textures, Vaidyanathan et al., Siggraph 2023

But which way is “correct”?

NVIDIA

But which way is "correct"?

Bold question:
Are we teaching and using texture filtering *"incorrectly"*?!

NVIDIA

Literature review and historical precedents

Precedent: Pre-multiplied alpha

filt(target)*filt(alpha) + (1-filt(alpha))*source

"Compositing digital images", Thomas Porter and Tom Duff., SIGGRAPH 1984.
Figure credit: “premultiplied alpha – 2022”, Inigo Quilez

Precedent: Pre-multiplied alpha

filt(target)*filt(alpha) + (1-filt(alpha))*source

filt(target*alpha) + (1-filt(alpha))*source

"Compositing digital images", Thomas Porter and Tom Duff., SIGGRAPH 1984.
Figure credit: “premultiplied alpha – 2022”, Inigo Quilez

Precedent: Percentage Closer Shadow Filtering

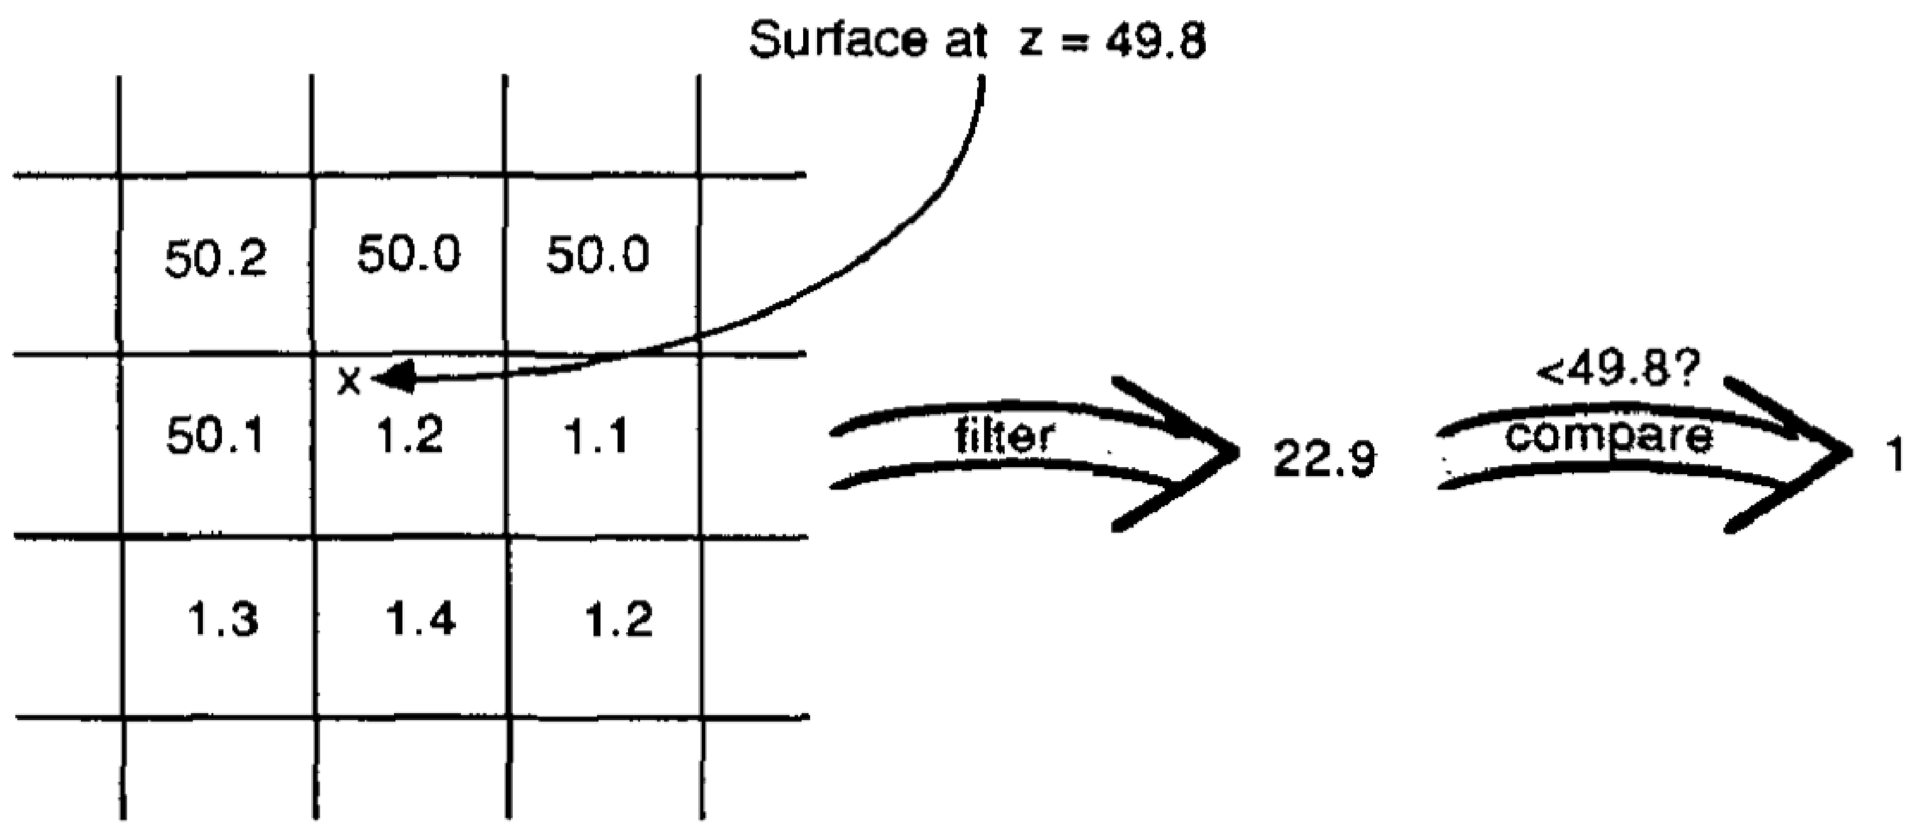

a) Ordinary texture map filtering. Does not work for depth maps.

$$visibility = z < \int depth(u, v)\, \mathrm{d}u\, \mathrm{d}v$$

Rendering Antialiased Shadows With Depth Maps, Reeves et al., SIGGRAPH 1987.

Precedent: Percentage Closer Shadow Filtering

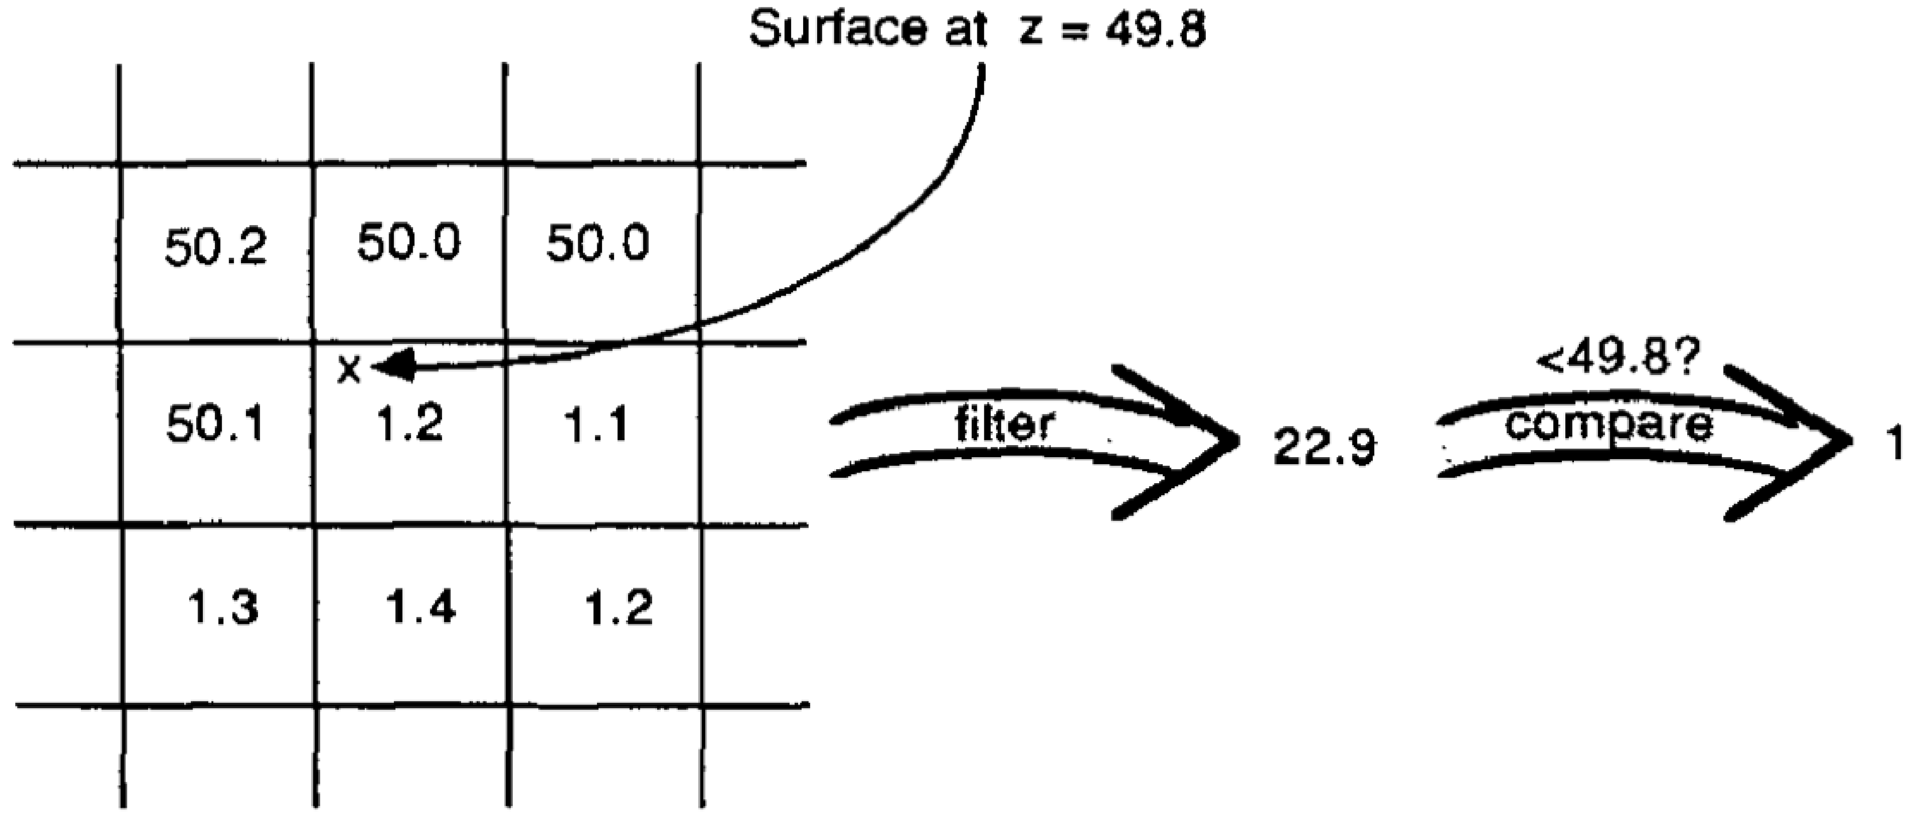

a) Ordinary texture map filtering. Does not work for depth maps.

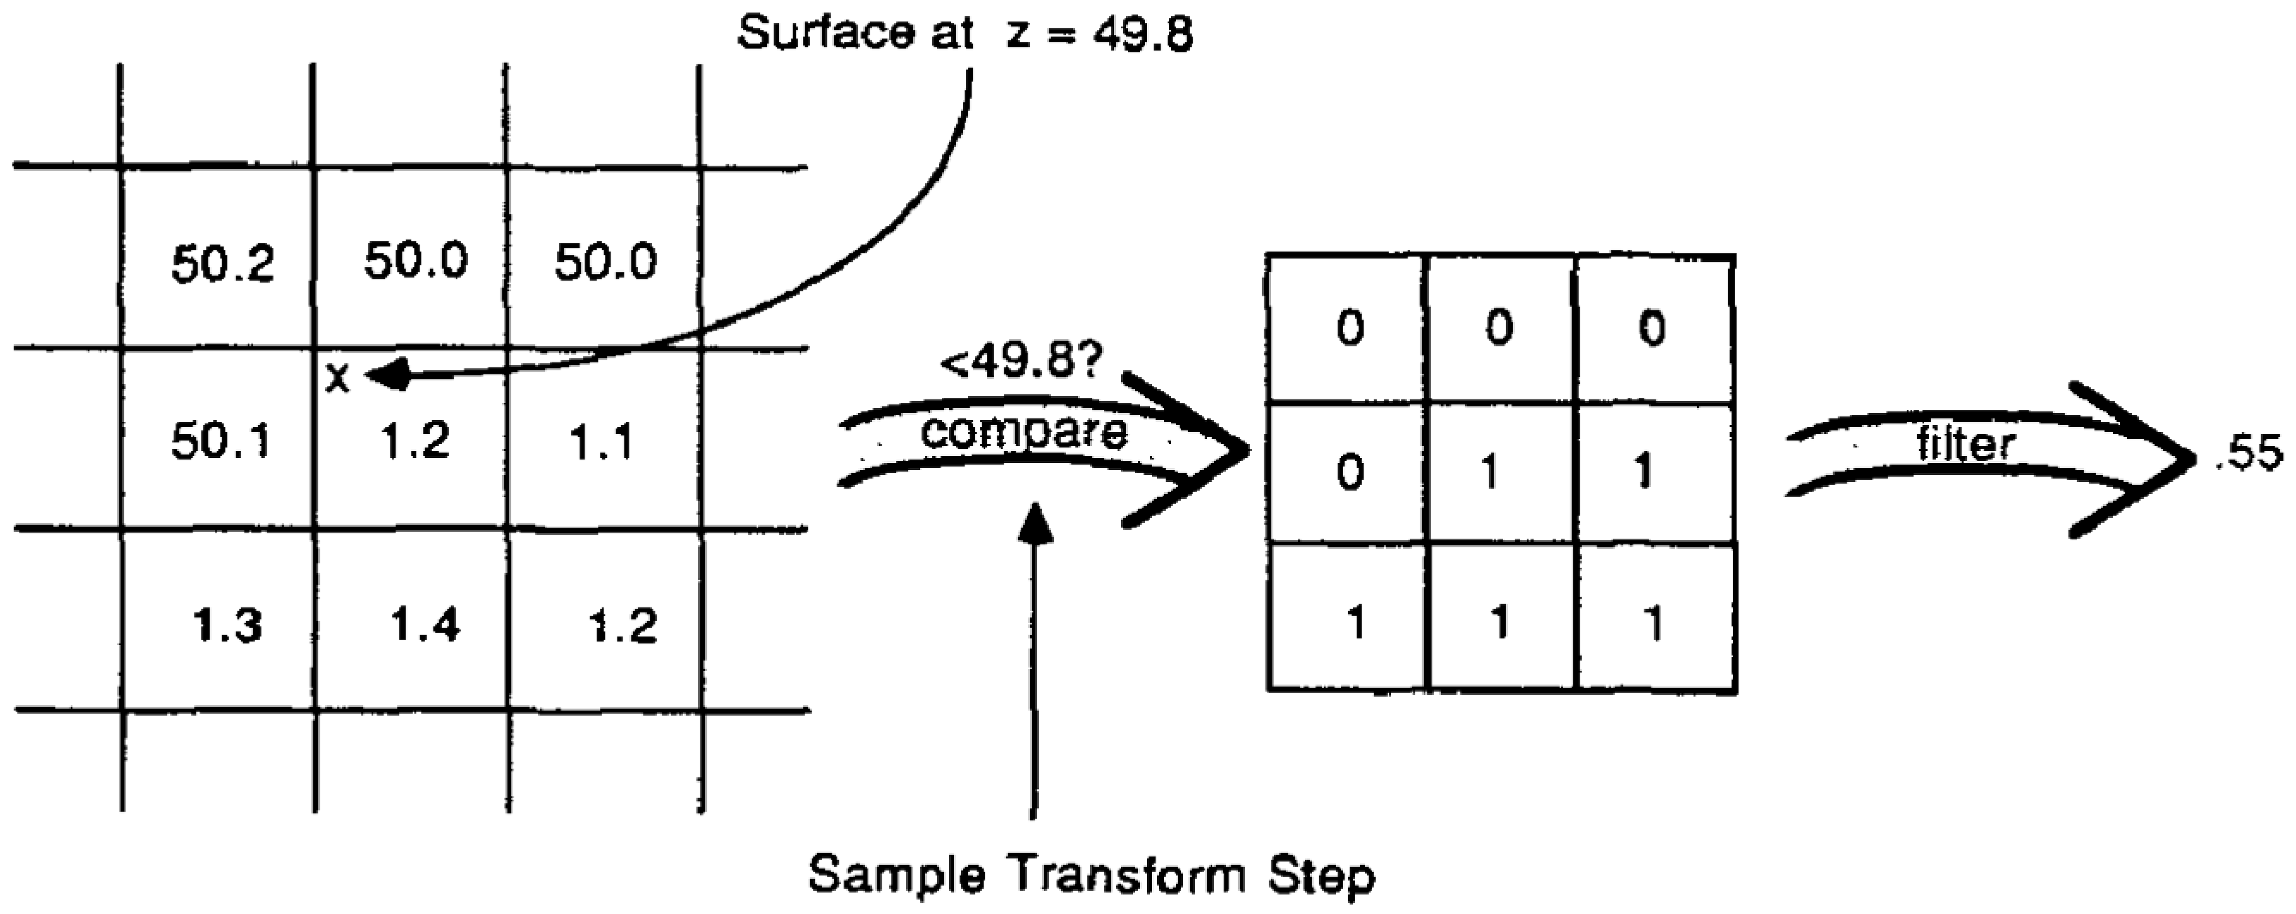

b) Percentage closer filtering.

$$visibility = z < \int depth(u,v)\, \mathrm{d}u\, \mathrm{d}v$$

$$visibility = \int f(u,v)\,(z < depth(u,v))\, \mathrm{d}u\, \mathrm{d}v$$

Jitter (u,v) to sample f

***Rendering Antialiased Shadows With Depth Maps*, Reeves et al., SIGGRAPH 1987.**

Precedent: Percentage Closer Shadow Filtering

Finally, we hope to be able to generalize and formalize the sample transformation step in percentage closer filtering. We believe that this technique may have important implications to the use of texture maps for other purposes. For example, in bump mapping [Bli78], specular reflections could be computed before filtering, and the results could be filtered and sampled as ordinary textures. In this way, specular highlights from the microfacets of a bumpy surface would be maintained even as the surface were translated back into the far distance.

Rendering Antialiased Shadows With Depth Maps, Reeves et al., SIGGRAPH 1987.

Precedent: Specular anti-aliasing (minification)

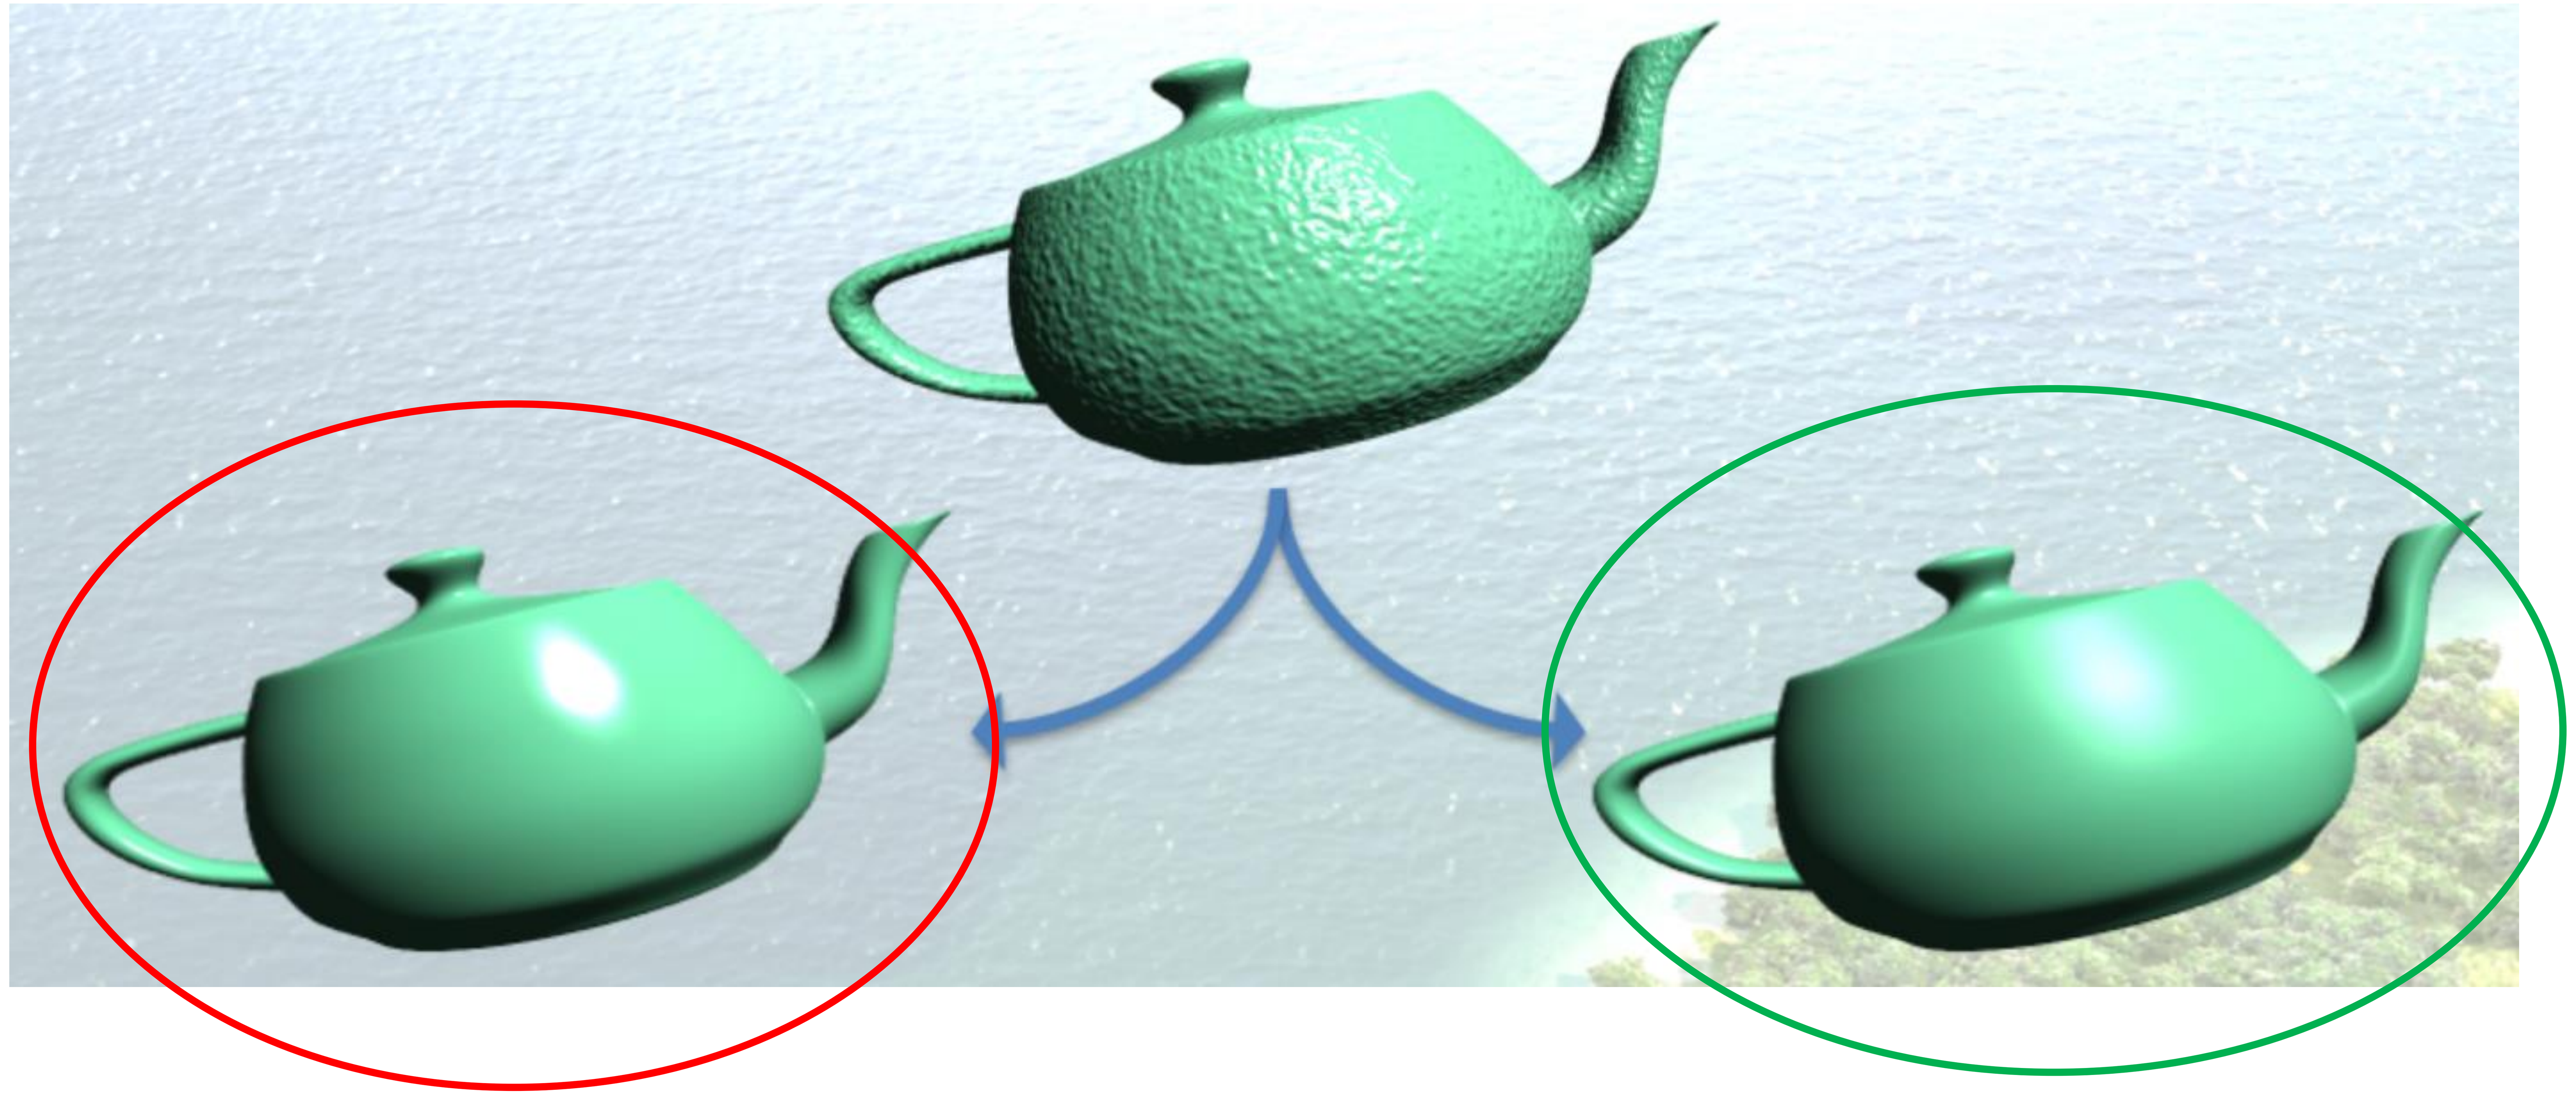

Mipmapping Normal Maps, Michael Toksvig, 2006
LEAN Mapping, Marc Olano and Dan Baker, I3D 2011
Figure credit*: Spectacular Specular: LEAN and CLEAN Specular Highlights*, Dan Baker, GDC 2011

Still a problem today: Metalness vs specular PBR workflow

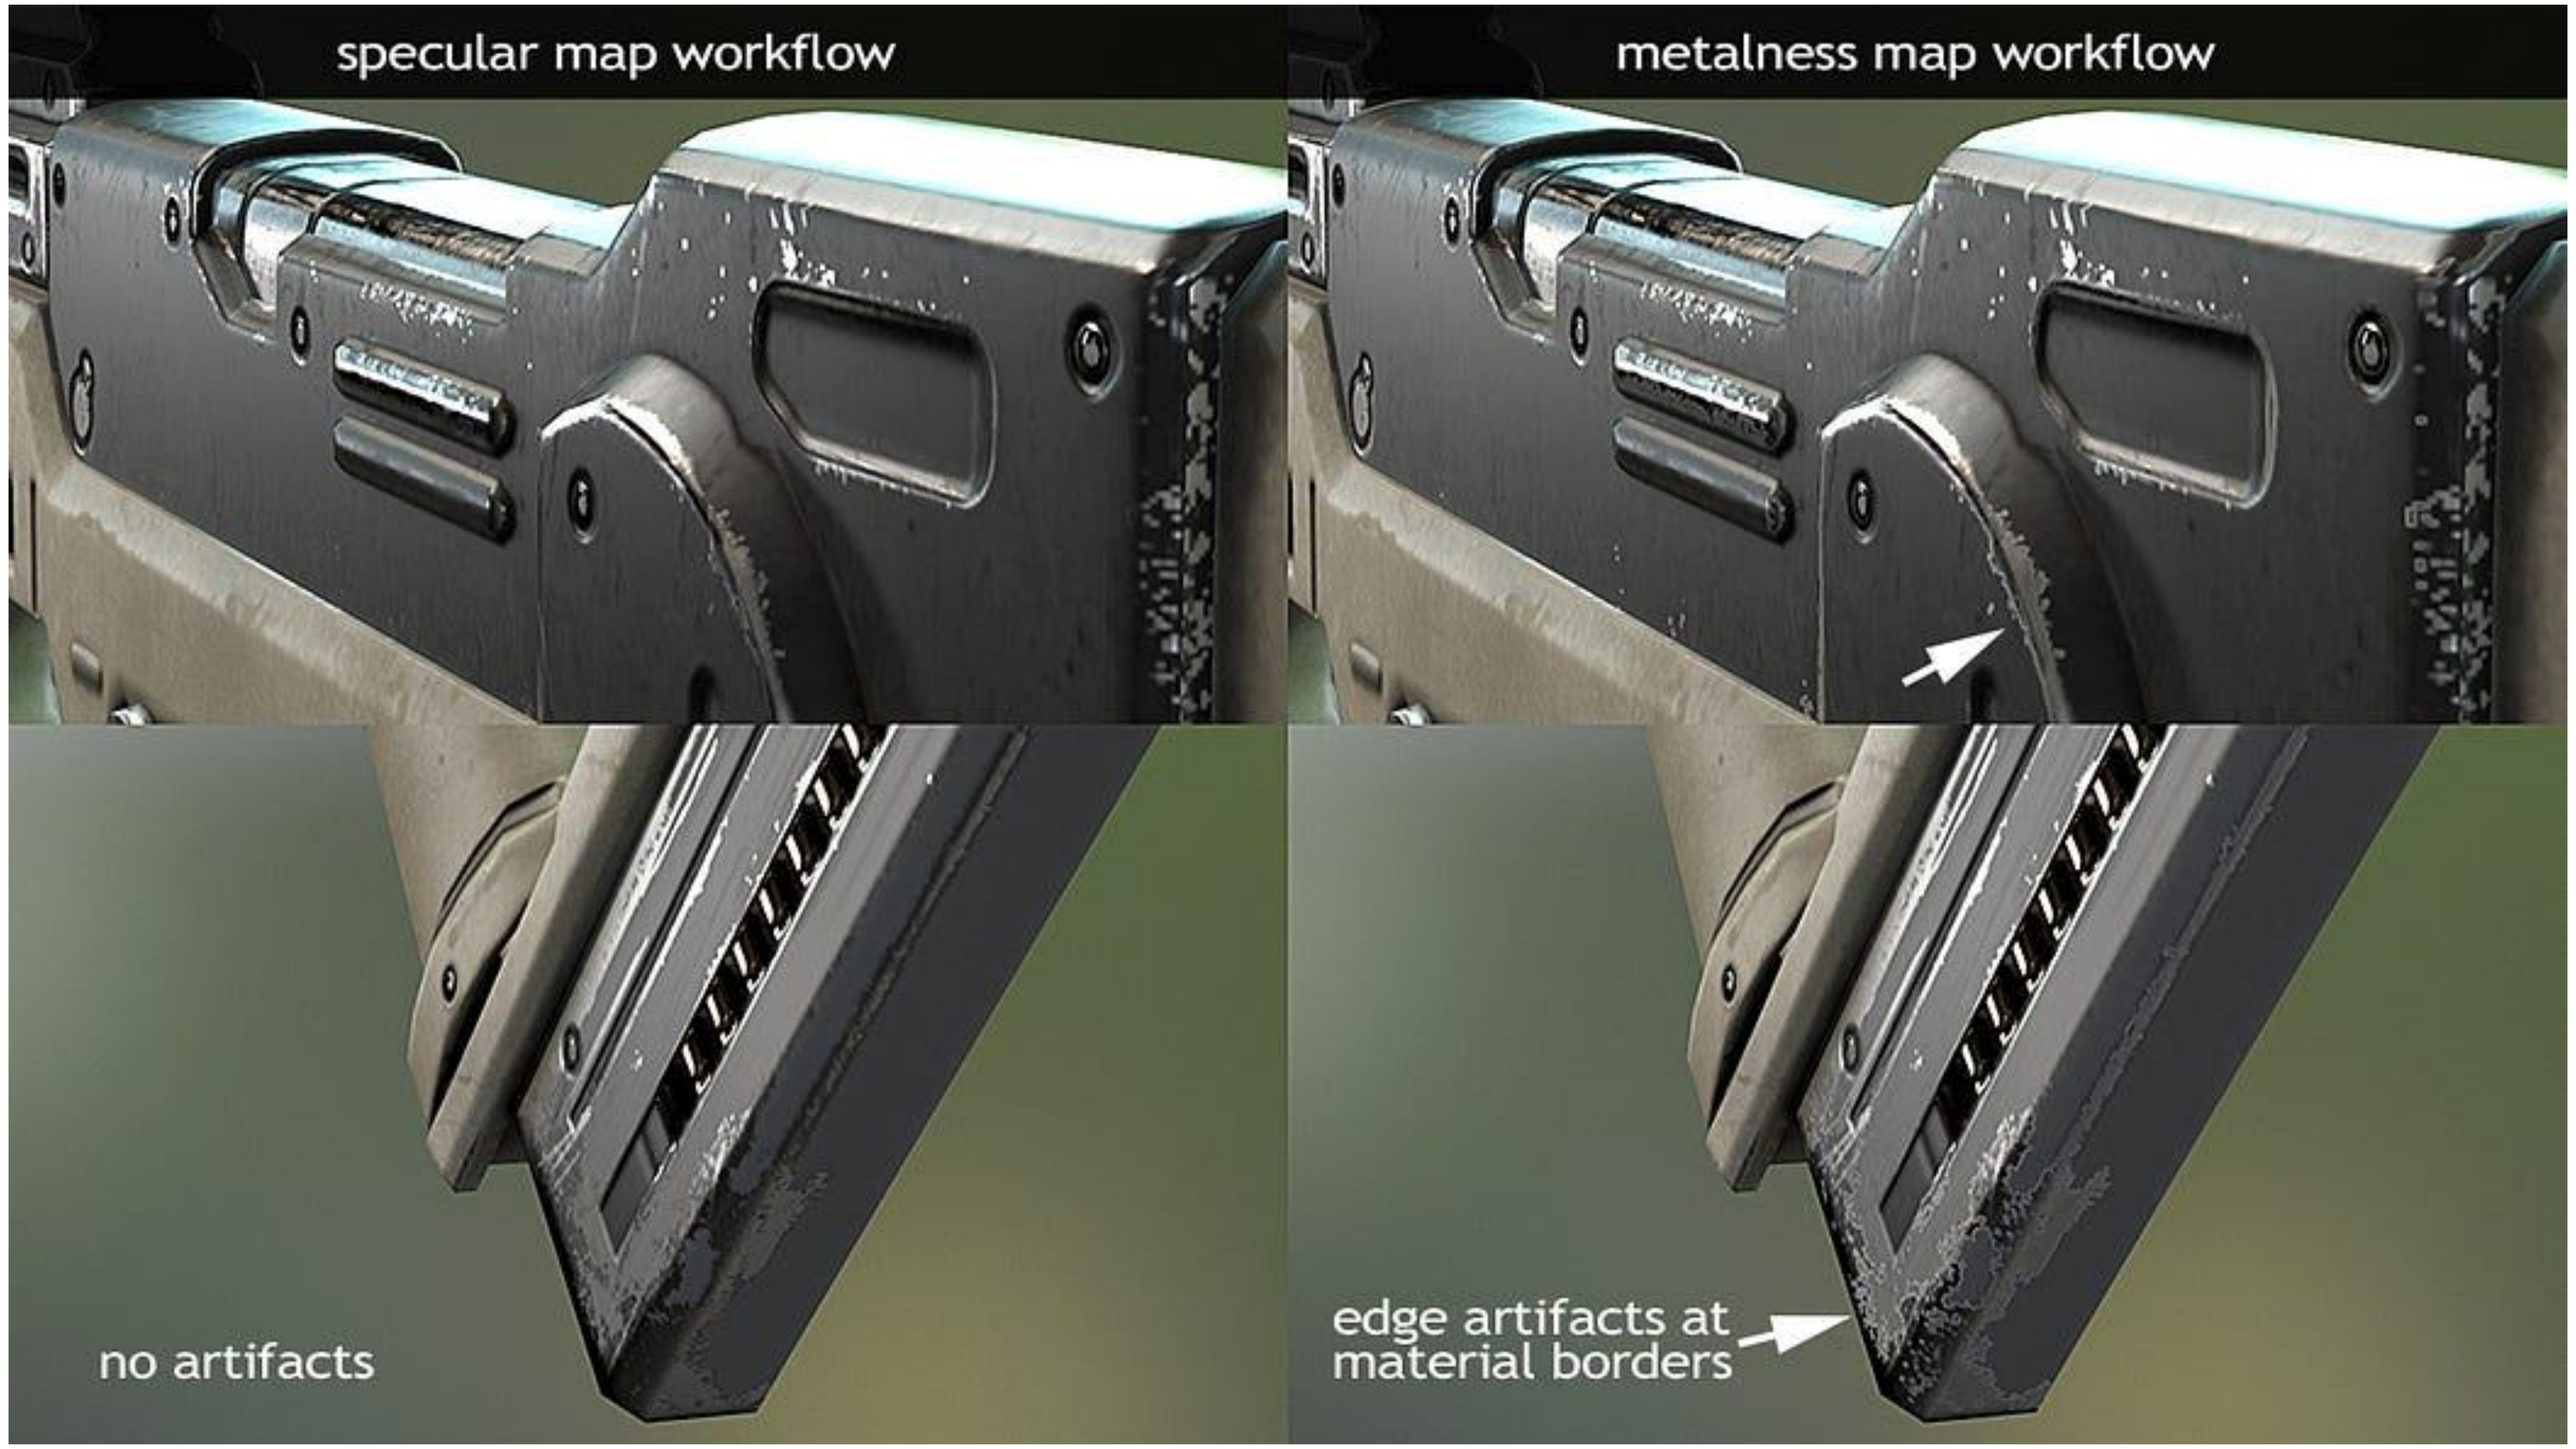

Specular reflectance = lerp(0.04, filt(color), filt(metalness))
Diffuse reflectance = lerp(filt(color), 0.0, filt(metalness))

Figure credit*: Metallic magic,* Daniel Rose

Motivation - summary

- Texture filtering theory and practice were developed for interpolating just “color”
- …in early work, not even gamma-corrected!

Motivation - summary

- Texture filtering theory and practice were developed for interpolating just “color”
- …in early work, not even gamma-corrected!
- Filtering and affine functions commute perfectly – this approach didn’t introduce errors
- **Non-linearity and filtering do not commute and swapping the order results in error**

Motivation - summary

- **Assumption**: textures are authored by artists with ~1-1 pixel-texel ratio

Motivation - summary

- **Assumption**: textures are authored by artists with ~1-1 pixel-texel ratio
- Minifying or magnifying textures before (non-linear) shading **introduces error/bias**
- Different techniques proposed to address specific types of errors

Motivation - summary

- Can we do better **in general**?
- Let's try to answer this question from 37y ago!

Finally, we hope to be able to generalize and formalize the sample transformation step in percentage closer filtering. We believe that this technique may have important implications to the use of texture maps for other purposes. For example, in bump mapping [Bli78], specular reflections could be computed before filtering, and the results could be filtered and sampled as ordinary textures. In this way, specular highlights from the microfacets of a bumpy surface would be maintained even as the surface were translated back into the far distance.

Rendering Antialiased Shadows With Depth Maps, Reeves et al., SIGGRAPH 1987.

Textures and the Rendering Equation

Filtering Before Shading (Standard Practice Today)

shading

$$L_o(p,\omega_o) = \overbrace{\int_{\mathbb{S}^2} f_r(\omega_i,\omega_o)\, L_i(p,\omega_i)\, |\cos\theta_i|\, \mathrm{d}\omega_i}^{\text{shading}}$$

Filtering Before Shading (Standard Practice Today)

shading

$$L_o(p,\omega_o) = \underbrace{\int_{\mathbb{S}^2} f_r(\omega_i,\omega_o)\, L_i(p,\omega_i)\, |\cos\theta_i|\, \mathrm{d}\omega_i}_{\text{shading}}$$

$$\textbf{reflectance} = \underbrace{\int f(u,v)\, tex_{\text{reflectance}}(u,v)\, \mathrm{d}u\, \mathrm{d}v}_{\text{filtering}}$$

$$\textbf{normal} = \underbrace{\int f(u,v)\, tex_{\text{normalmap}}(u,v)\, \mathrm{d}u\, \mathrm{d}v}_{\text{filtering}}$$

filtering

filtering

Filtering Before Shading (Standard Practice Today)

$$L_o(p,\omega_o) = \overbrace{\int_{\mathbb{S}^2} f_r(\omega_i,\omega_o)\, L_i(p,\omega_i)\, |\cos\theta_i|\, \mathrm{d}\omega_i}^{\text{shading}}$$

$$\textbf{reflectance} = \underbrace{\int f(u,v)\, tex_{\text{reflectance}}(u,v)\, \mathrm{d}u\, \mathrm{d}v}_{\text{filtering}}$$

$$\textbf{normal} = \underbrace{\int f(u,v)\, tex_{\text{normalmap}}(u,v)\, \mathrm{d}u\, \mathrm{d}v}_{\text{filtering}}$$

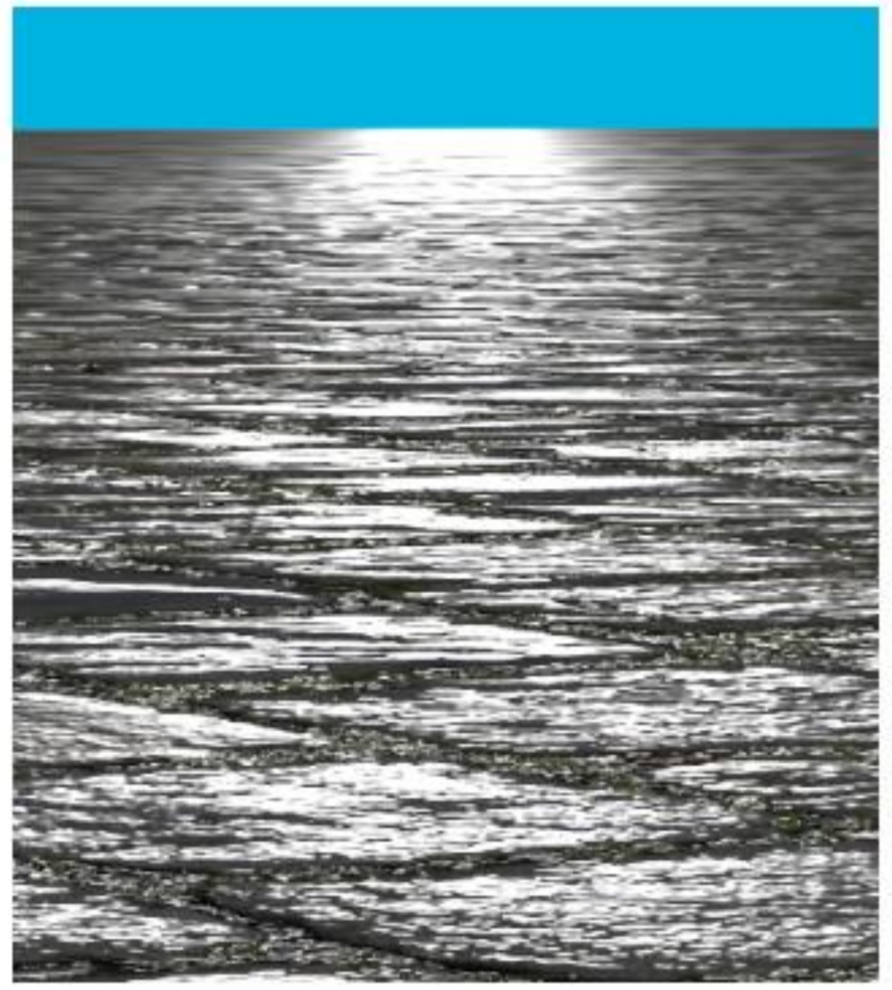

HW 16x Aniso

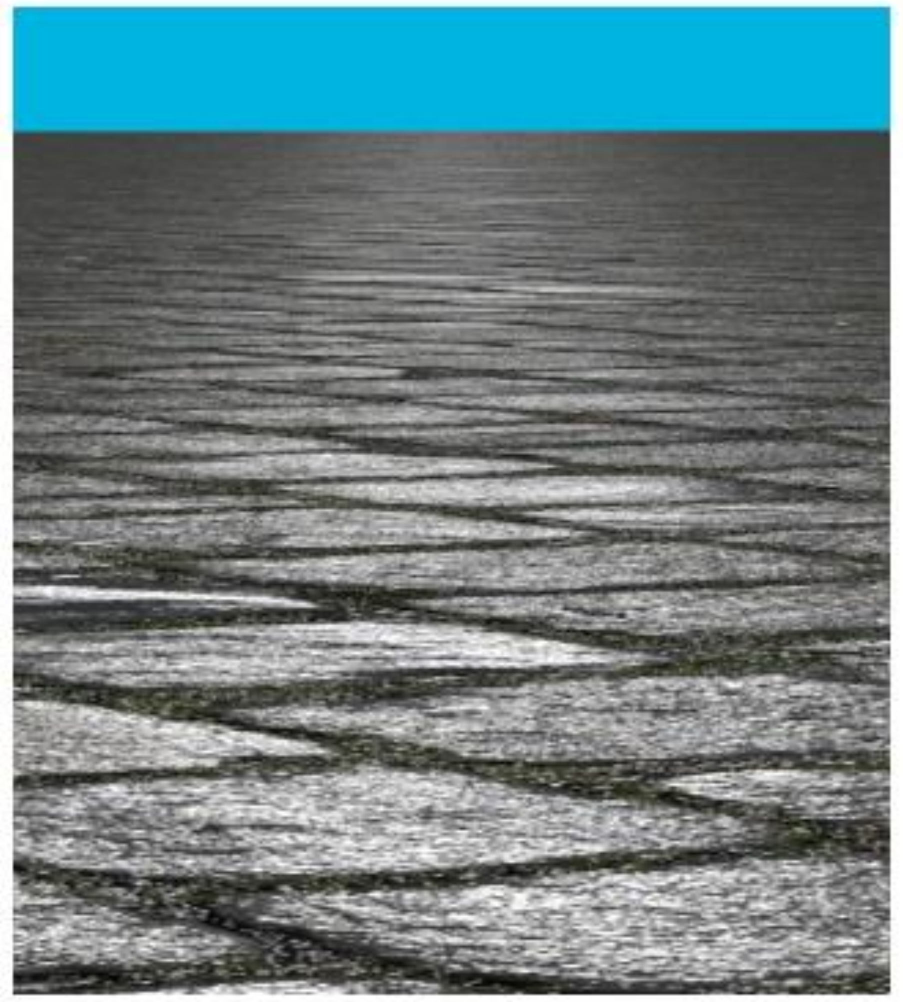

Reference

Filtering After Shading

$$L_o(p,\omega_o) = \overbrace{\int f(u,v)}^{\text{filtering}} \left(\overbrace{\int_{\mathbb{S}^2} f_r(\omega_i,\omega_o)\, L_i(p,\omega_i)\, |\cos\theta_i|\, \mathrm{d}\omega_i}^{\text{shading}} \right) \mathrm{d}u\, \mathrm{d}v$$

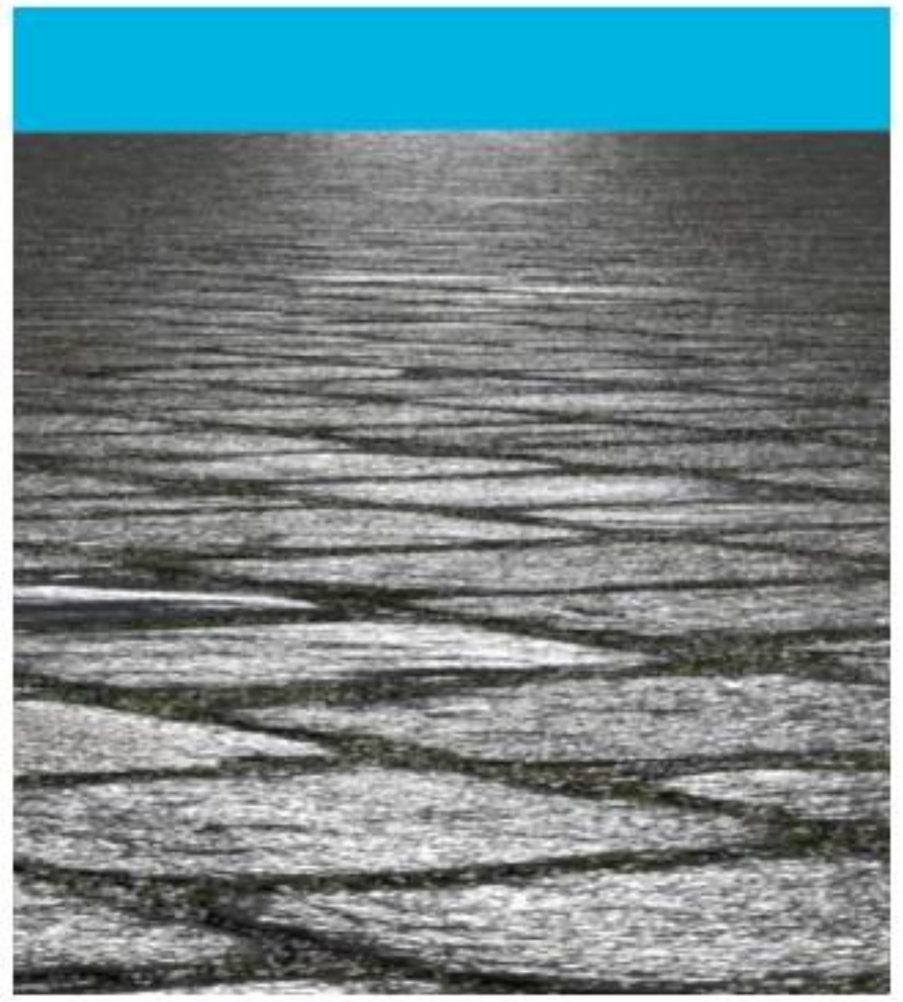

Filter after shading
(real-time implementation)

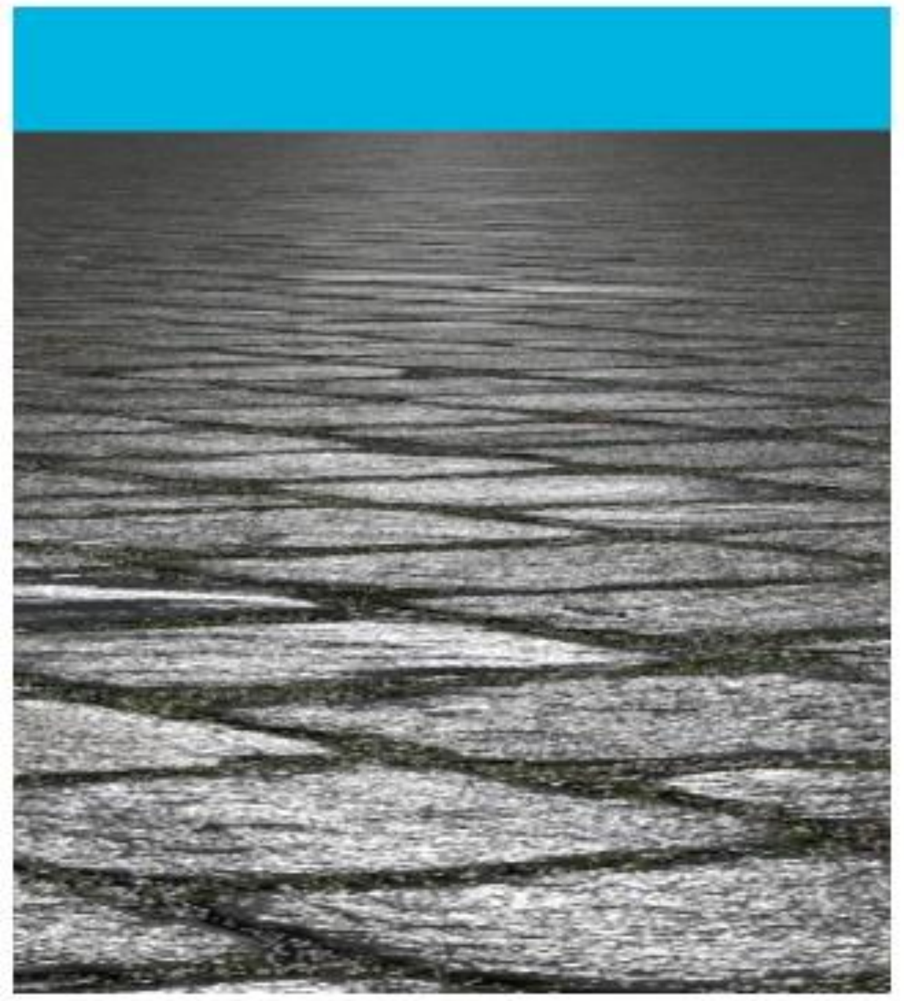

Reference

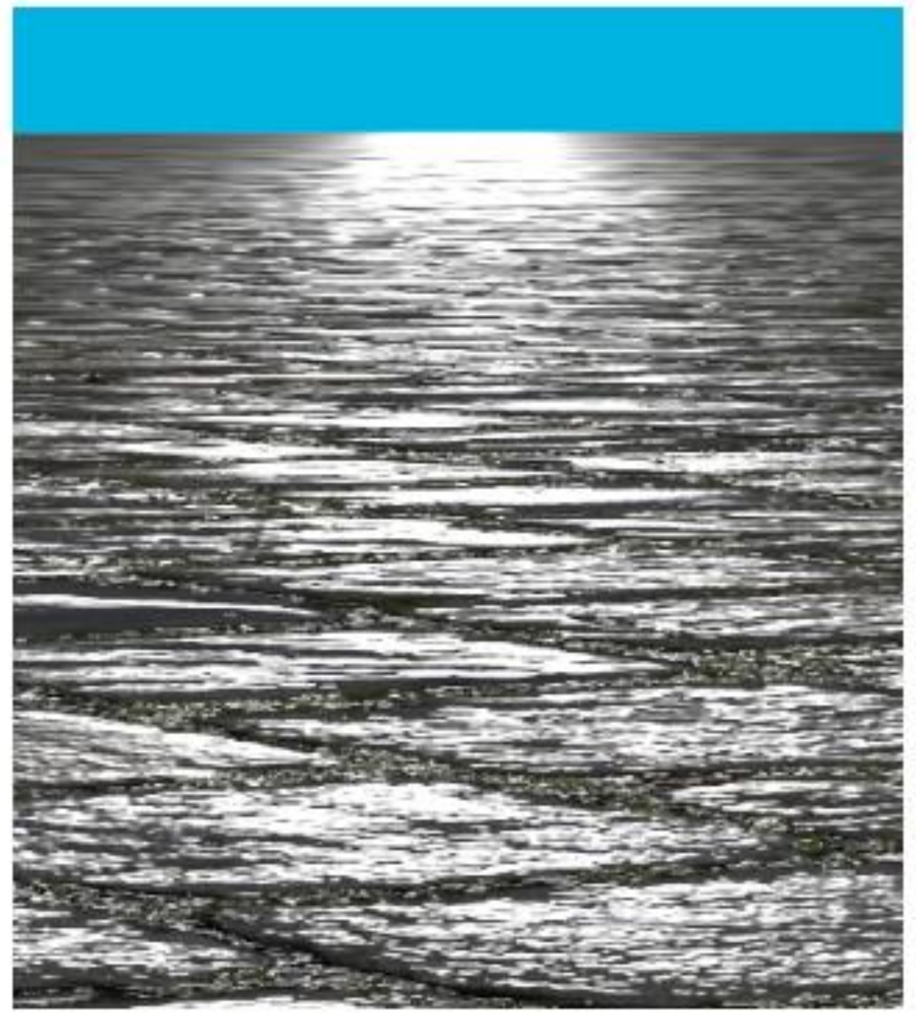

Filter before shading

Filtering After Shading

$$L_o(p, \omega_o) = \overbrace{\int f(u, v)}^{\text{filtering}} \overbrace{\left(\int_{\mathbb{S}^2} f_r(\omega_i, \omega_o)\, L_i(p, \omega_i)\, |\cos\theta_i|\, \mathrm{d}\omega_i \right)}^{\text{shading}} \mathrm{d}u\, \mathrm{d}v$$

????

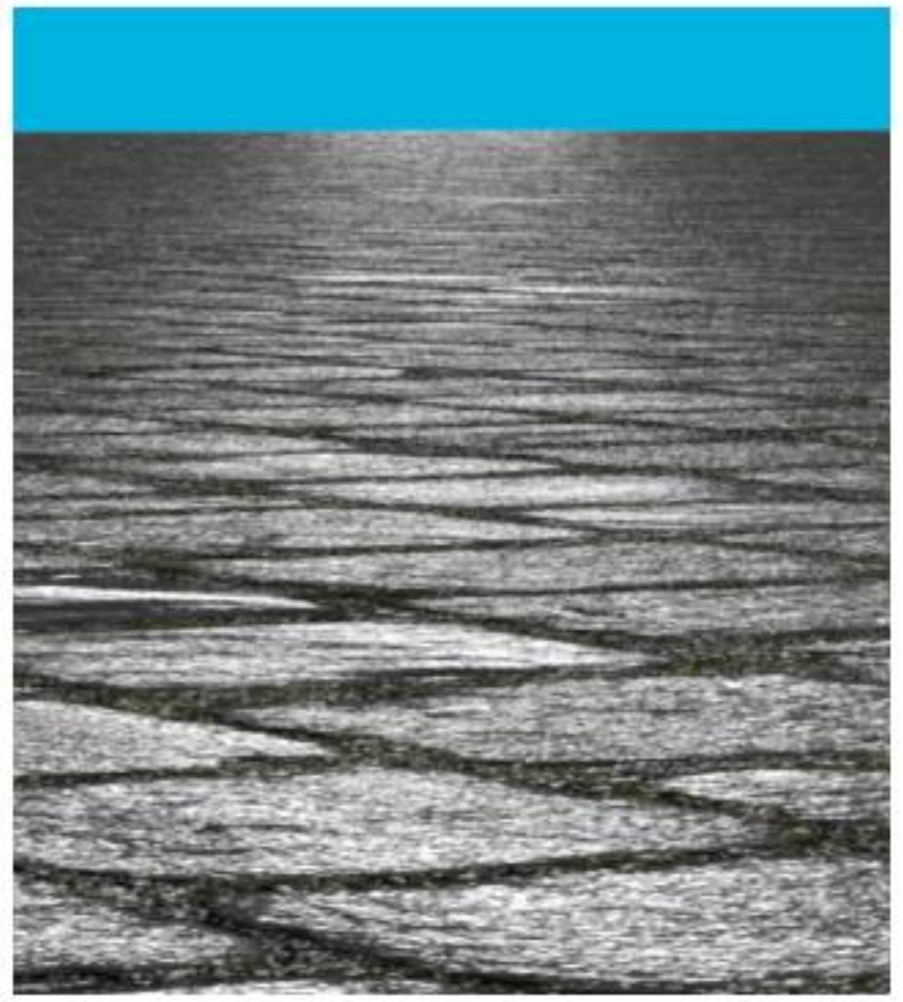

Filter after shading
(real-time implementation)

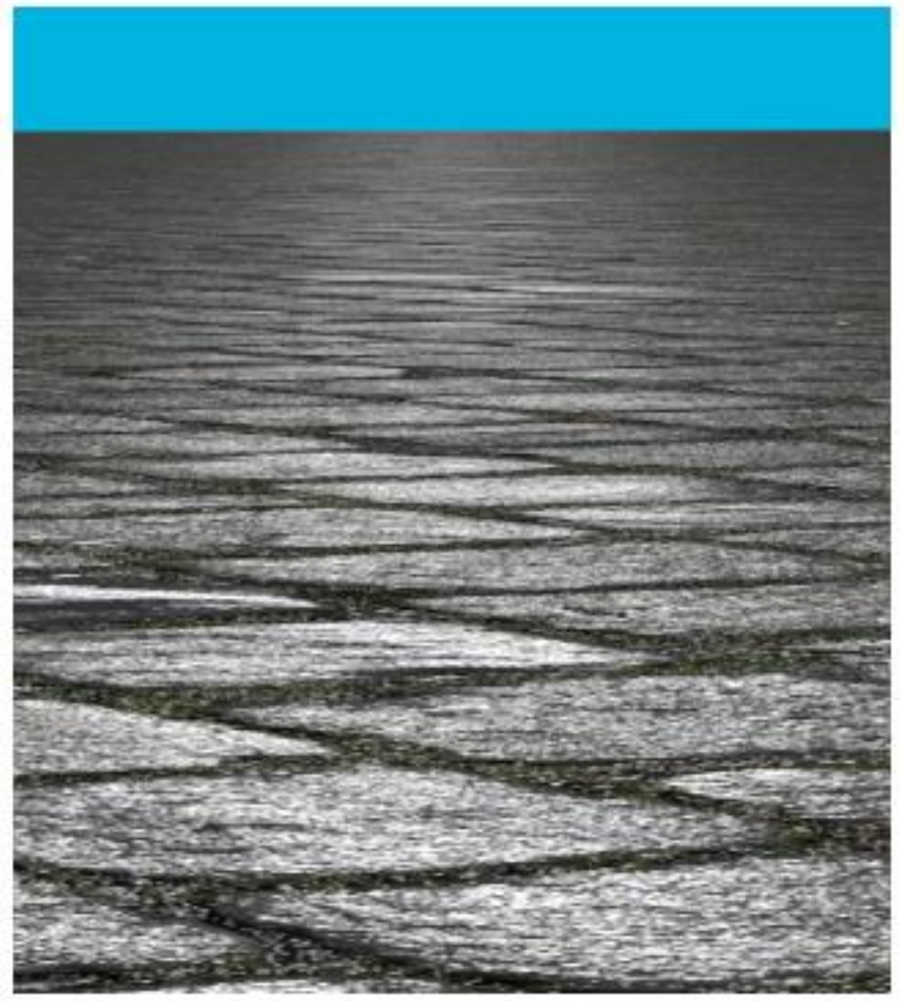

Reference

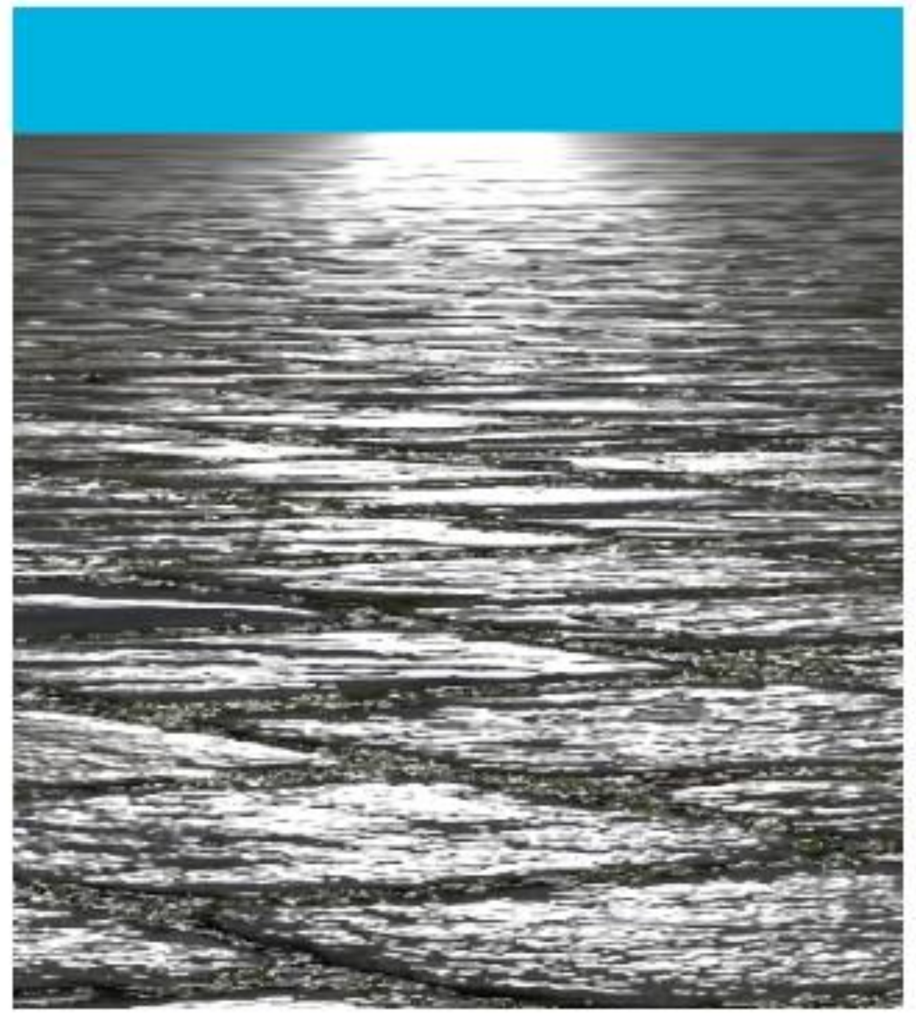

Filter before shading

Filtering After Shading

$$L_o(p,\omega_o) = \overbrace{\int f(u,v)}^{\text{filtering}} \left(\overbrace{\int_{\mathbb{S}^2} f_r(\omega_i,\omega_o)\, L_i(p,\omega_i)\, |\cos\theta_i|\, \mathrm{d}\omega_i}^{\text{shading}} \right) \mathrm{d}u\, \mathrm{d}v$$

- Use Monte Carlo!
- Sample $(u', v') \sim f(u,v)$

Filtering After Shading

$$L_o(p,\omega_o) = \overbrace{\int f(u,v)}^{\text{filtering}} \left(\overbrace{\int_{\mathbb{S}^2} f_r(\omega_i,\omega_o)\, L_i(p,\omega_i)\, |\cos\theta_i|\, \mathrm{d}\omega_i}^{\text{shading}} \right) \mathrm{d}u\, \mathrm{d}v$$

- Use Monte Carlo!
- Sample $(u',v') \sim f(u,v)$
- Estimator: $L_o(p,\omega_o) \approx \dfrac{f(u,v)}{p(u,v)} \left(\displaystyle\int_{\mathbb{S}^2} f_r(\omega_i,\omega_o)\, L_i(p,\omega_i)\, |\cos\theta_i|\, \mathrm{d}\omega_i \right) \mathrm{d}u\, \mathrm{d}v$

$$= \int_{\mathbb{S}^2} f_r(\omega_i,\omega_o)\, L_i(p,\omega_i)\, |\cos\theta_i|\, \mathrm{d}\omega_i$$

$\textbf{reflectance} = tex_{\text{reflectance}}(u',v')$

$\textbf{normal} = tex_{\text{normalmap}}(u',v')$

Unfiltered single texel lookups!

Even Single Sample!
Real Time Rendering – Noise

White noise

Single frame

Even Single Sample!
Real Time Rendering – Noise

- DLSS as the robust temporal integrator

White noise

Single frame

DLSS

Even Single Sample!
Real Time Rendering – Noise

- Spatiotemporal Blue Noise reduces the noise appearance

White noise

(Spatiotemporal) Blue Noise

Single frame

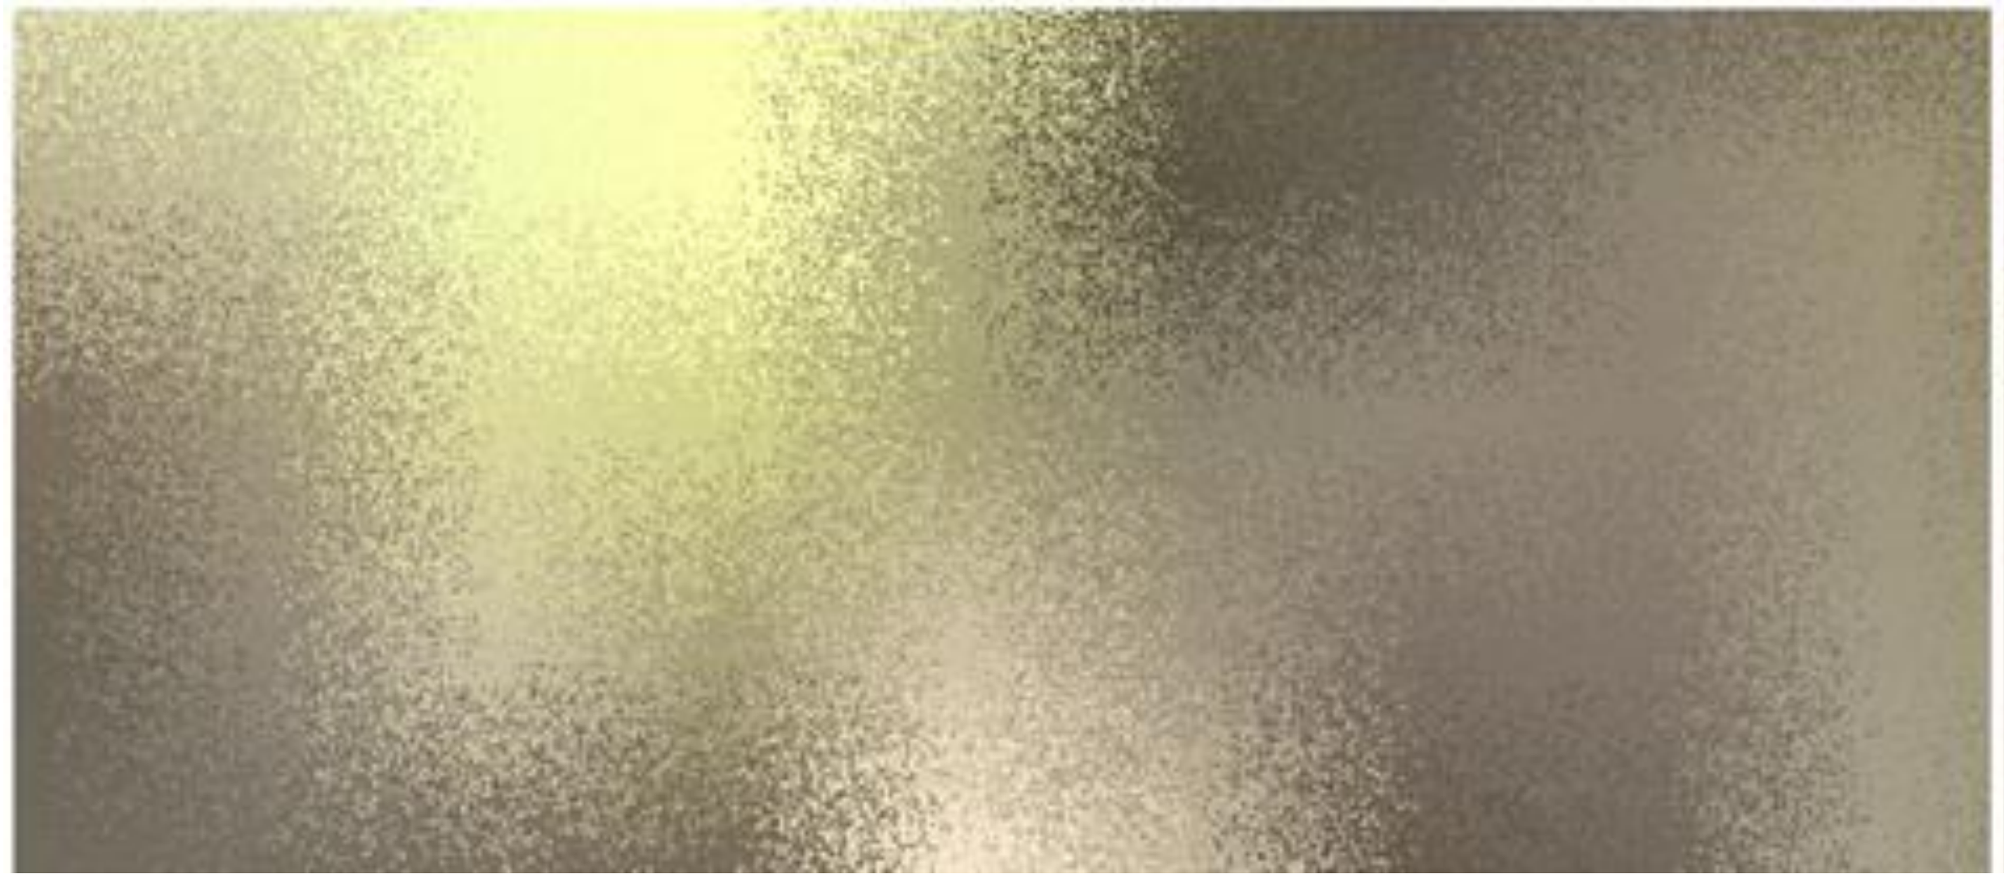

Even Single Sample!
Real Time Rendering – Noise

- DLSS as the robust temporal integrator
- Spatiotemporal Blue Noise reduces the noise appearance
- Makes DLSS job easier, together -> no visible noise in most cases

White noise | (Spatiotemporal) Blue Noise

Single frame

DLSS

Two families of methods

Sampling Texture Filters – Discrete, 1D

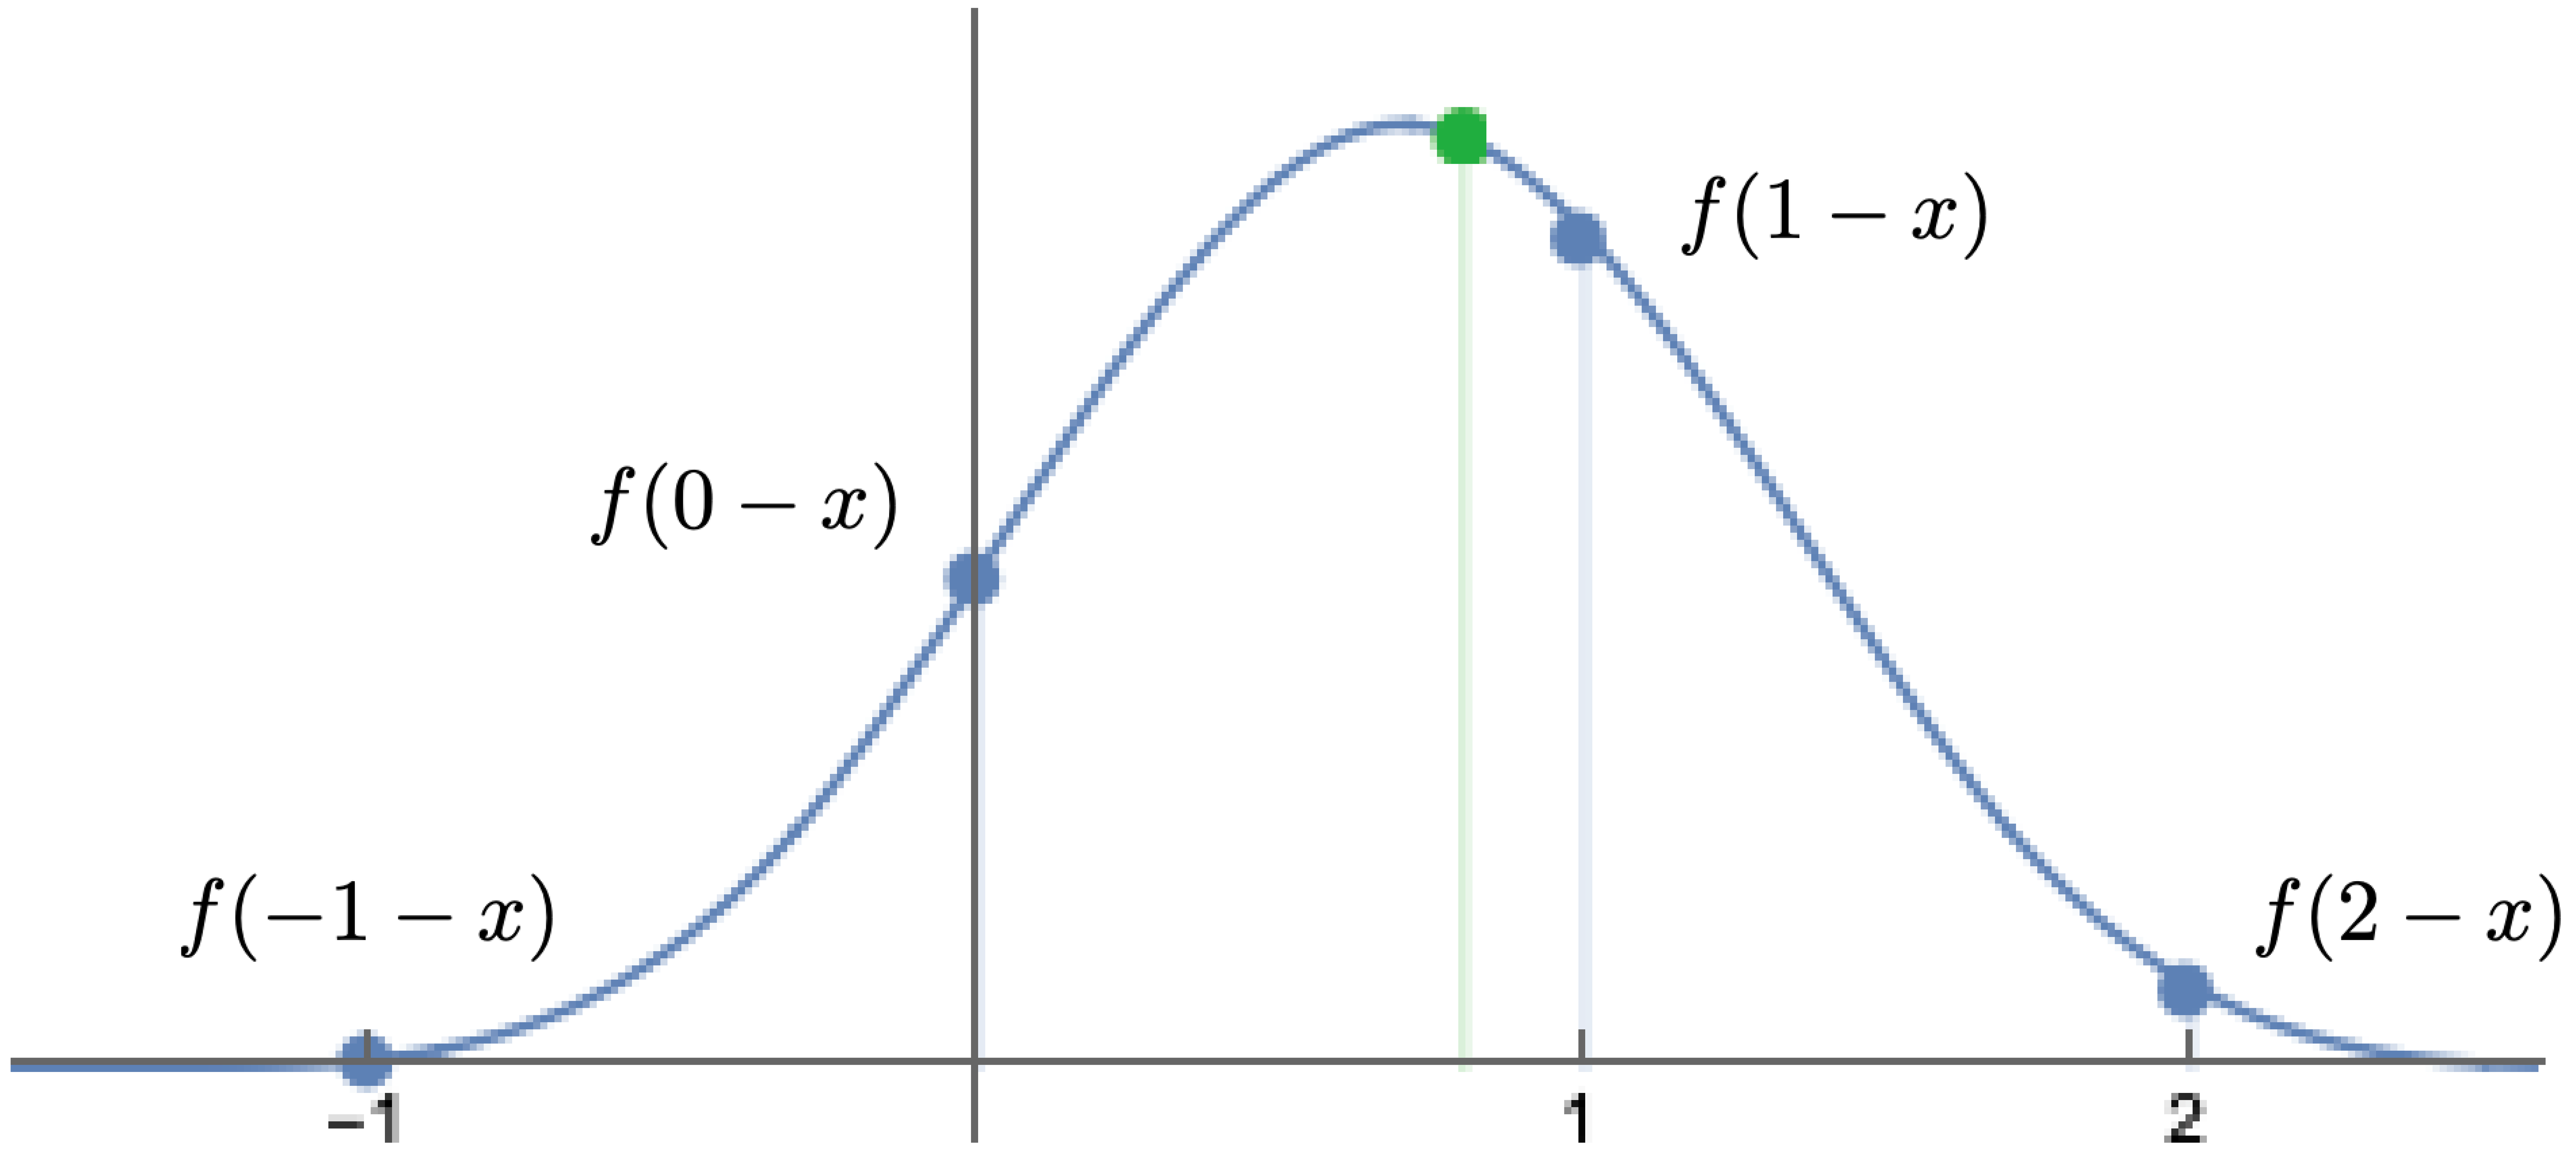

$$lookup(x) = \int f(u-x)\, t(u)\, \mathrm{d}u$$

Sampling Texture Filters – Discrete, 1D

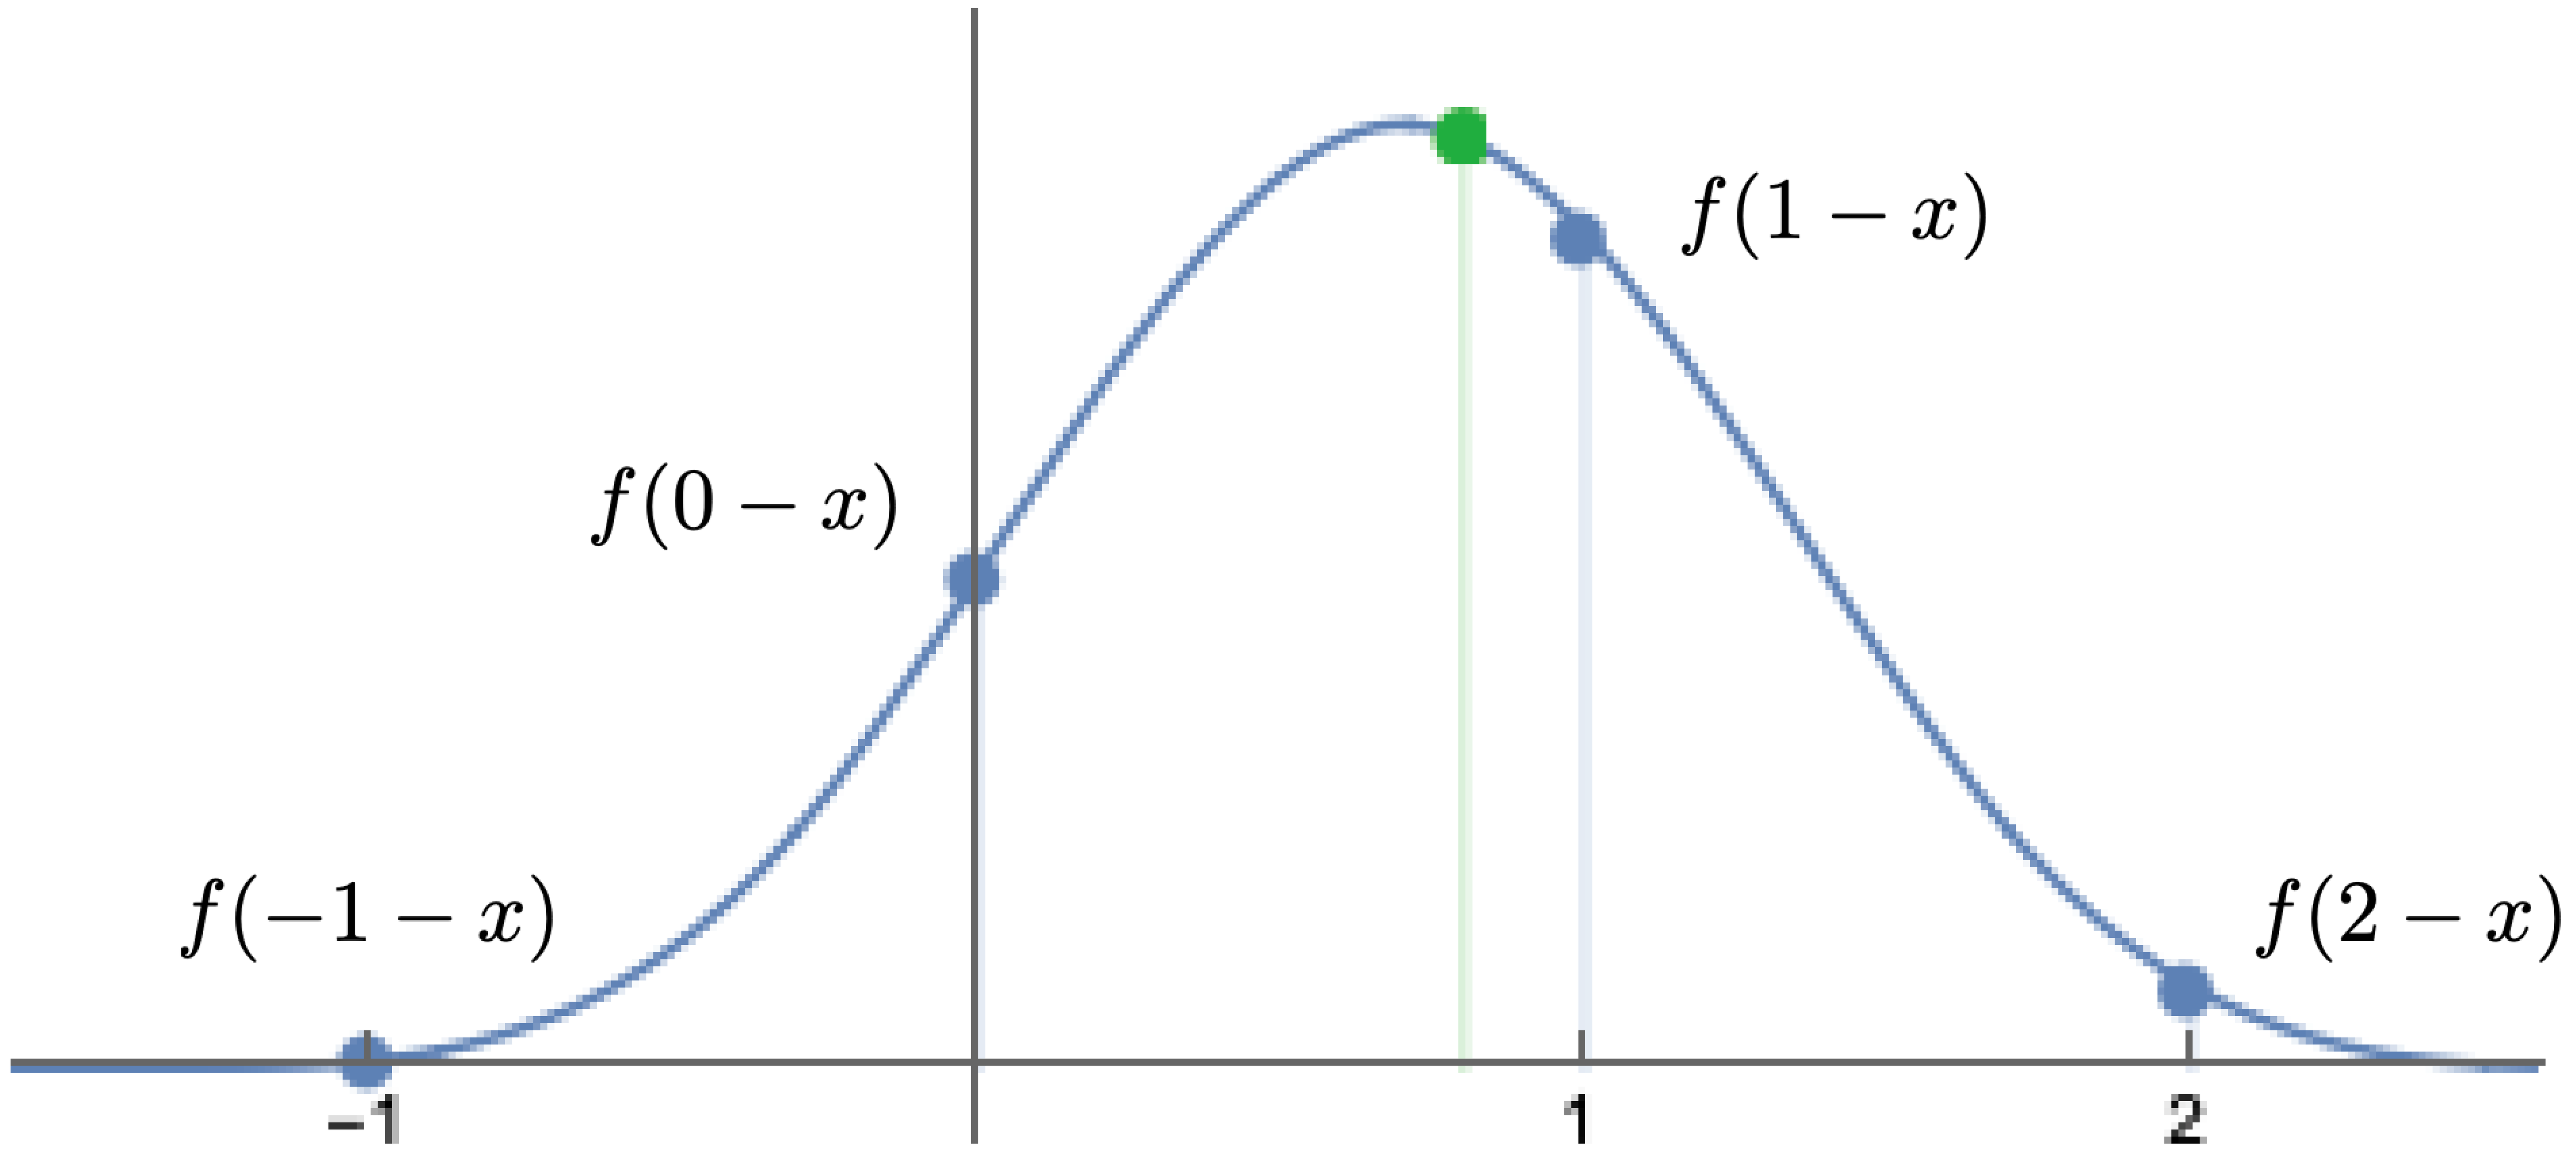

$$lookup(x) = \int f(u-x)\,t(u)\,\mathrm{d}u = \sum_{u=-1}^{2} f(u-x)\,t(u)$$

Sampling Texture Filters – Discrete, 1D

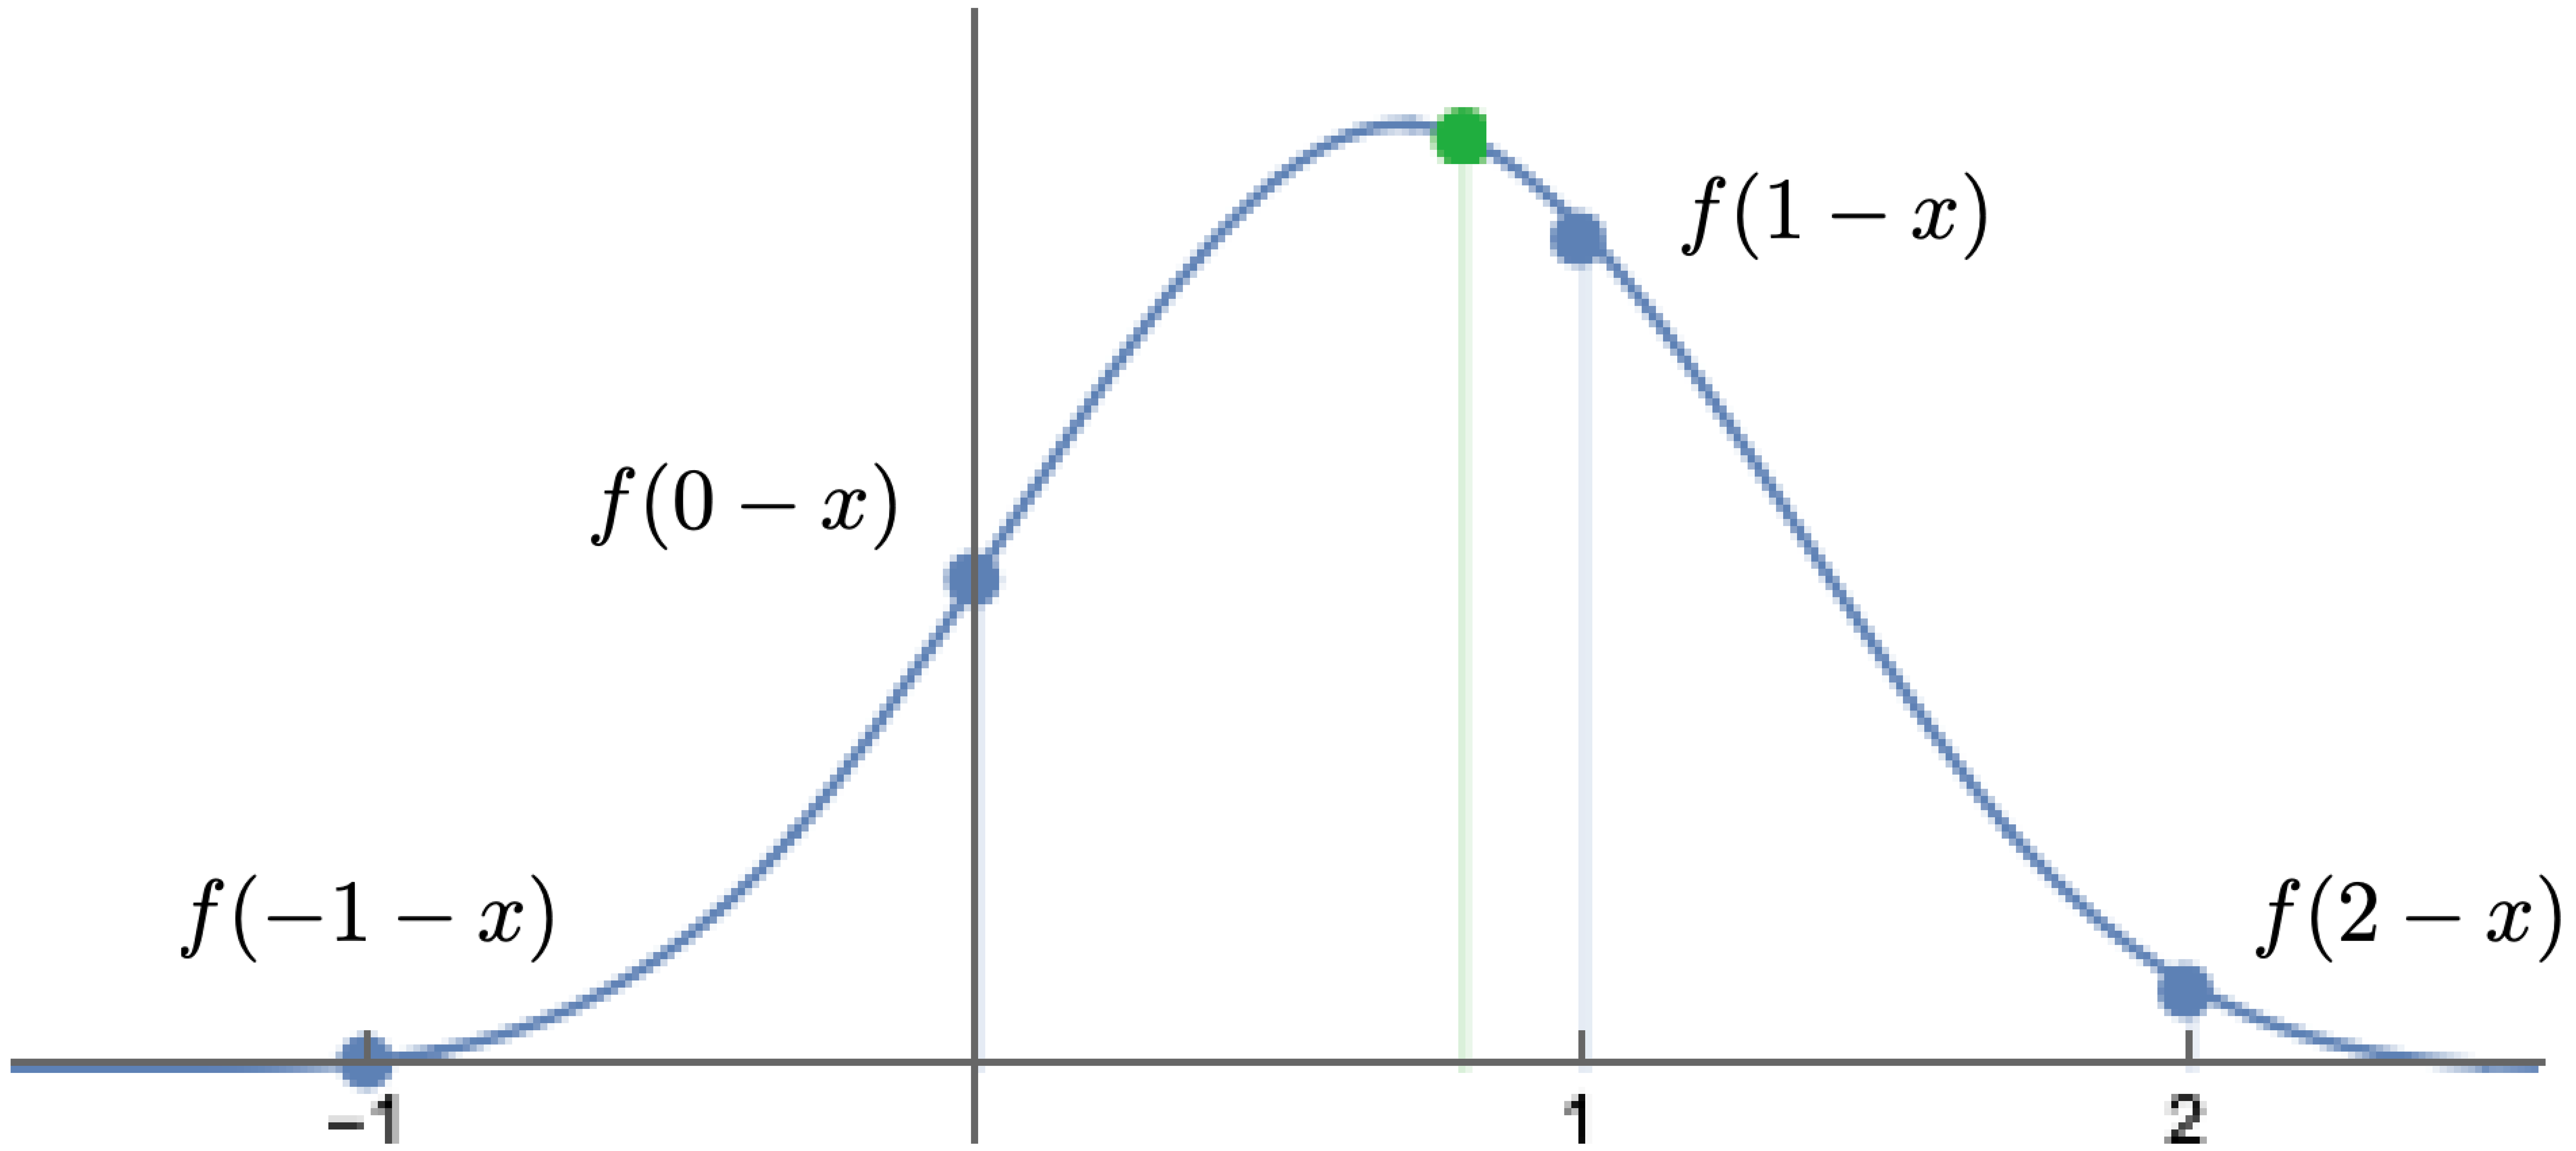

$$lookup(x) = \int f(u-x)\,t(u)\,\mathrm{d}u = \sum_{u=-1}^{2} \underbrace{f(u-x)}\,t(u)$$

Chose a sample with probability ~f

Filter Reservoir Sampling

- **Importance sampling**: Sample a texel with probability p~f
- Optimal if we don't know the signal and cannot sample the product
- Sample an array of weights or online through weighted reservoir sampling

Filter Reservoir Sampling

- **Importance sampling**: Sample a texel with probability p~f
- Optimal if we don't know the signal and cannot sample the product
- Sample an array of weights or online through weighted reservoir sampling

Multidimensional interpolating/approximating filters are mostly separable:

- Sample each dimension independently
- In d dimensions, the filtering cost for an n-tap filter is n*d, not n^d!

Sampling Texture Filters
Disadvantages of Filter Reservoir Sampling

- Discrete filter sampling – with large filters, can be costly
- Evaluate filter function K^M or K*M times
- Does not support infinite filters (Gaussian, sinc)

Sampling Texture Filters
Disadvantages of Filter Reservoir Sampling

- Discrete filter sampling – with large filters, can be costly
- Evaluate filter function K^M or K*M times
- Does not support infinite filters (Gaussian, sinc)

There's a different way!

- Let's analyze and understand the "UV jitter + nearest neighbor" prior work.

Magnification
What happens when you take a nearest-neighbor sample?

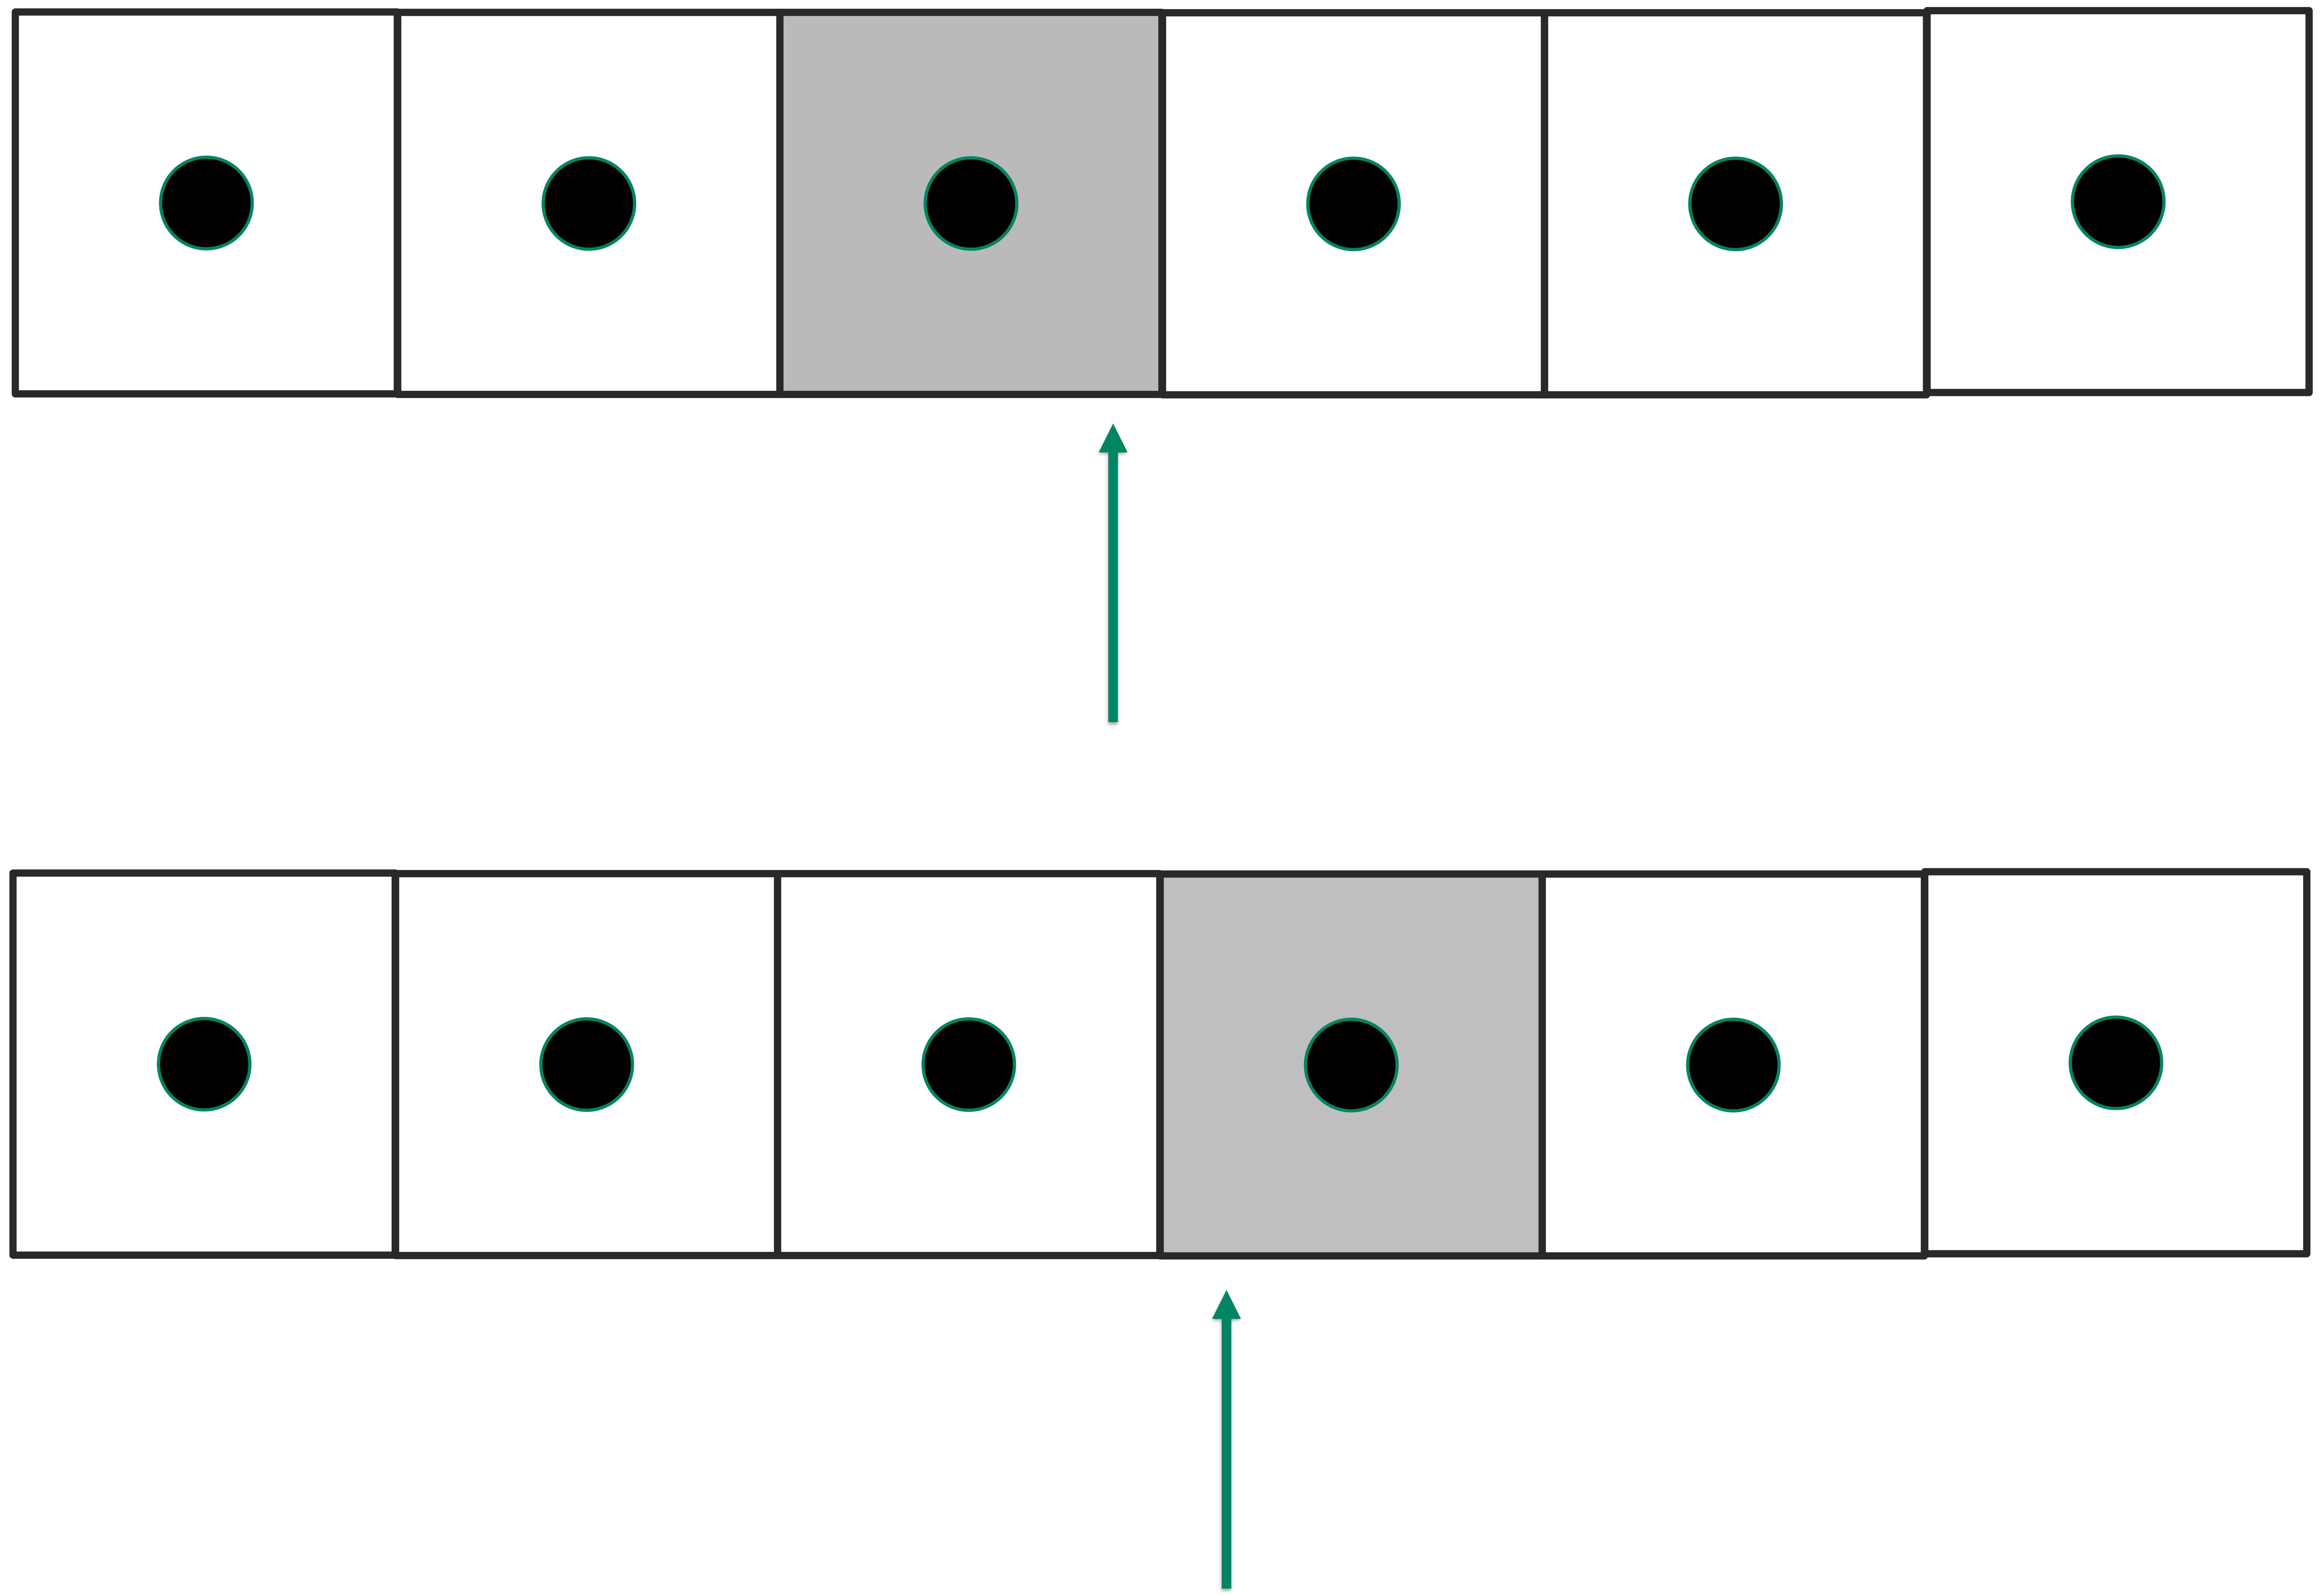

Magnification
What happens when you take a nearest-neighbor sample?

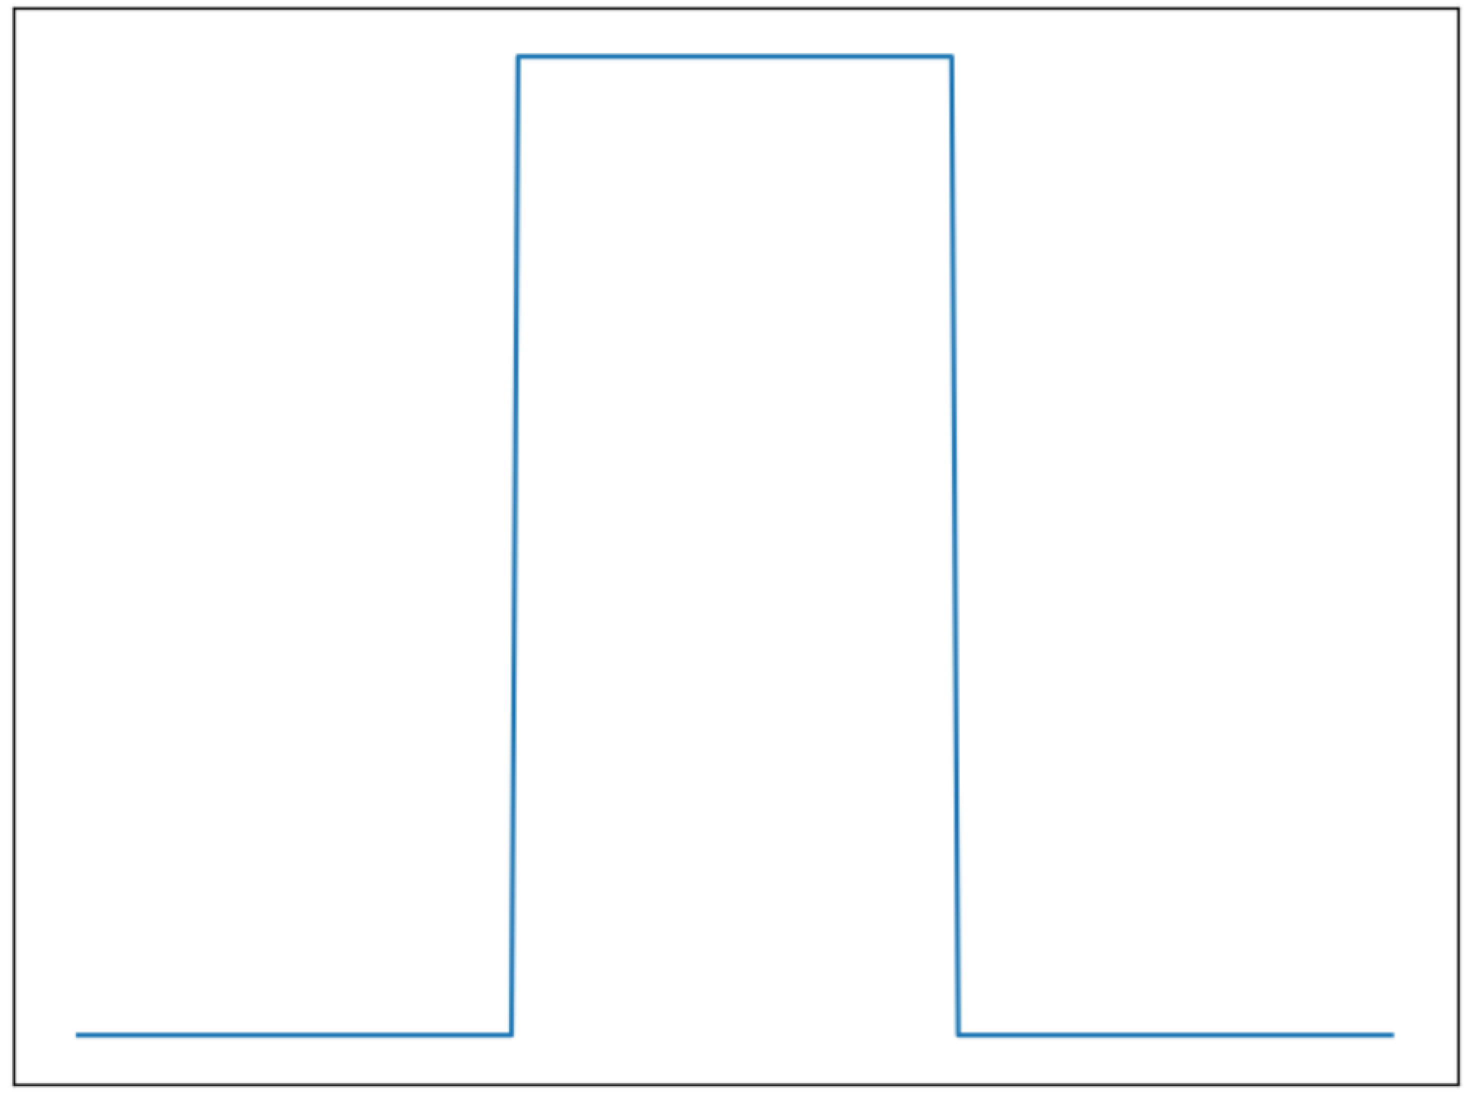

Nearest neighbor = box filter

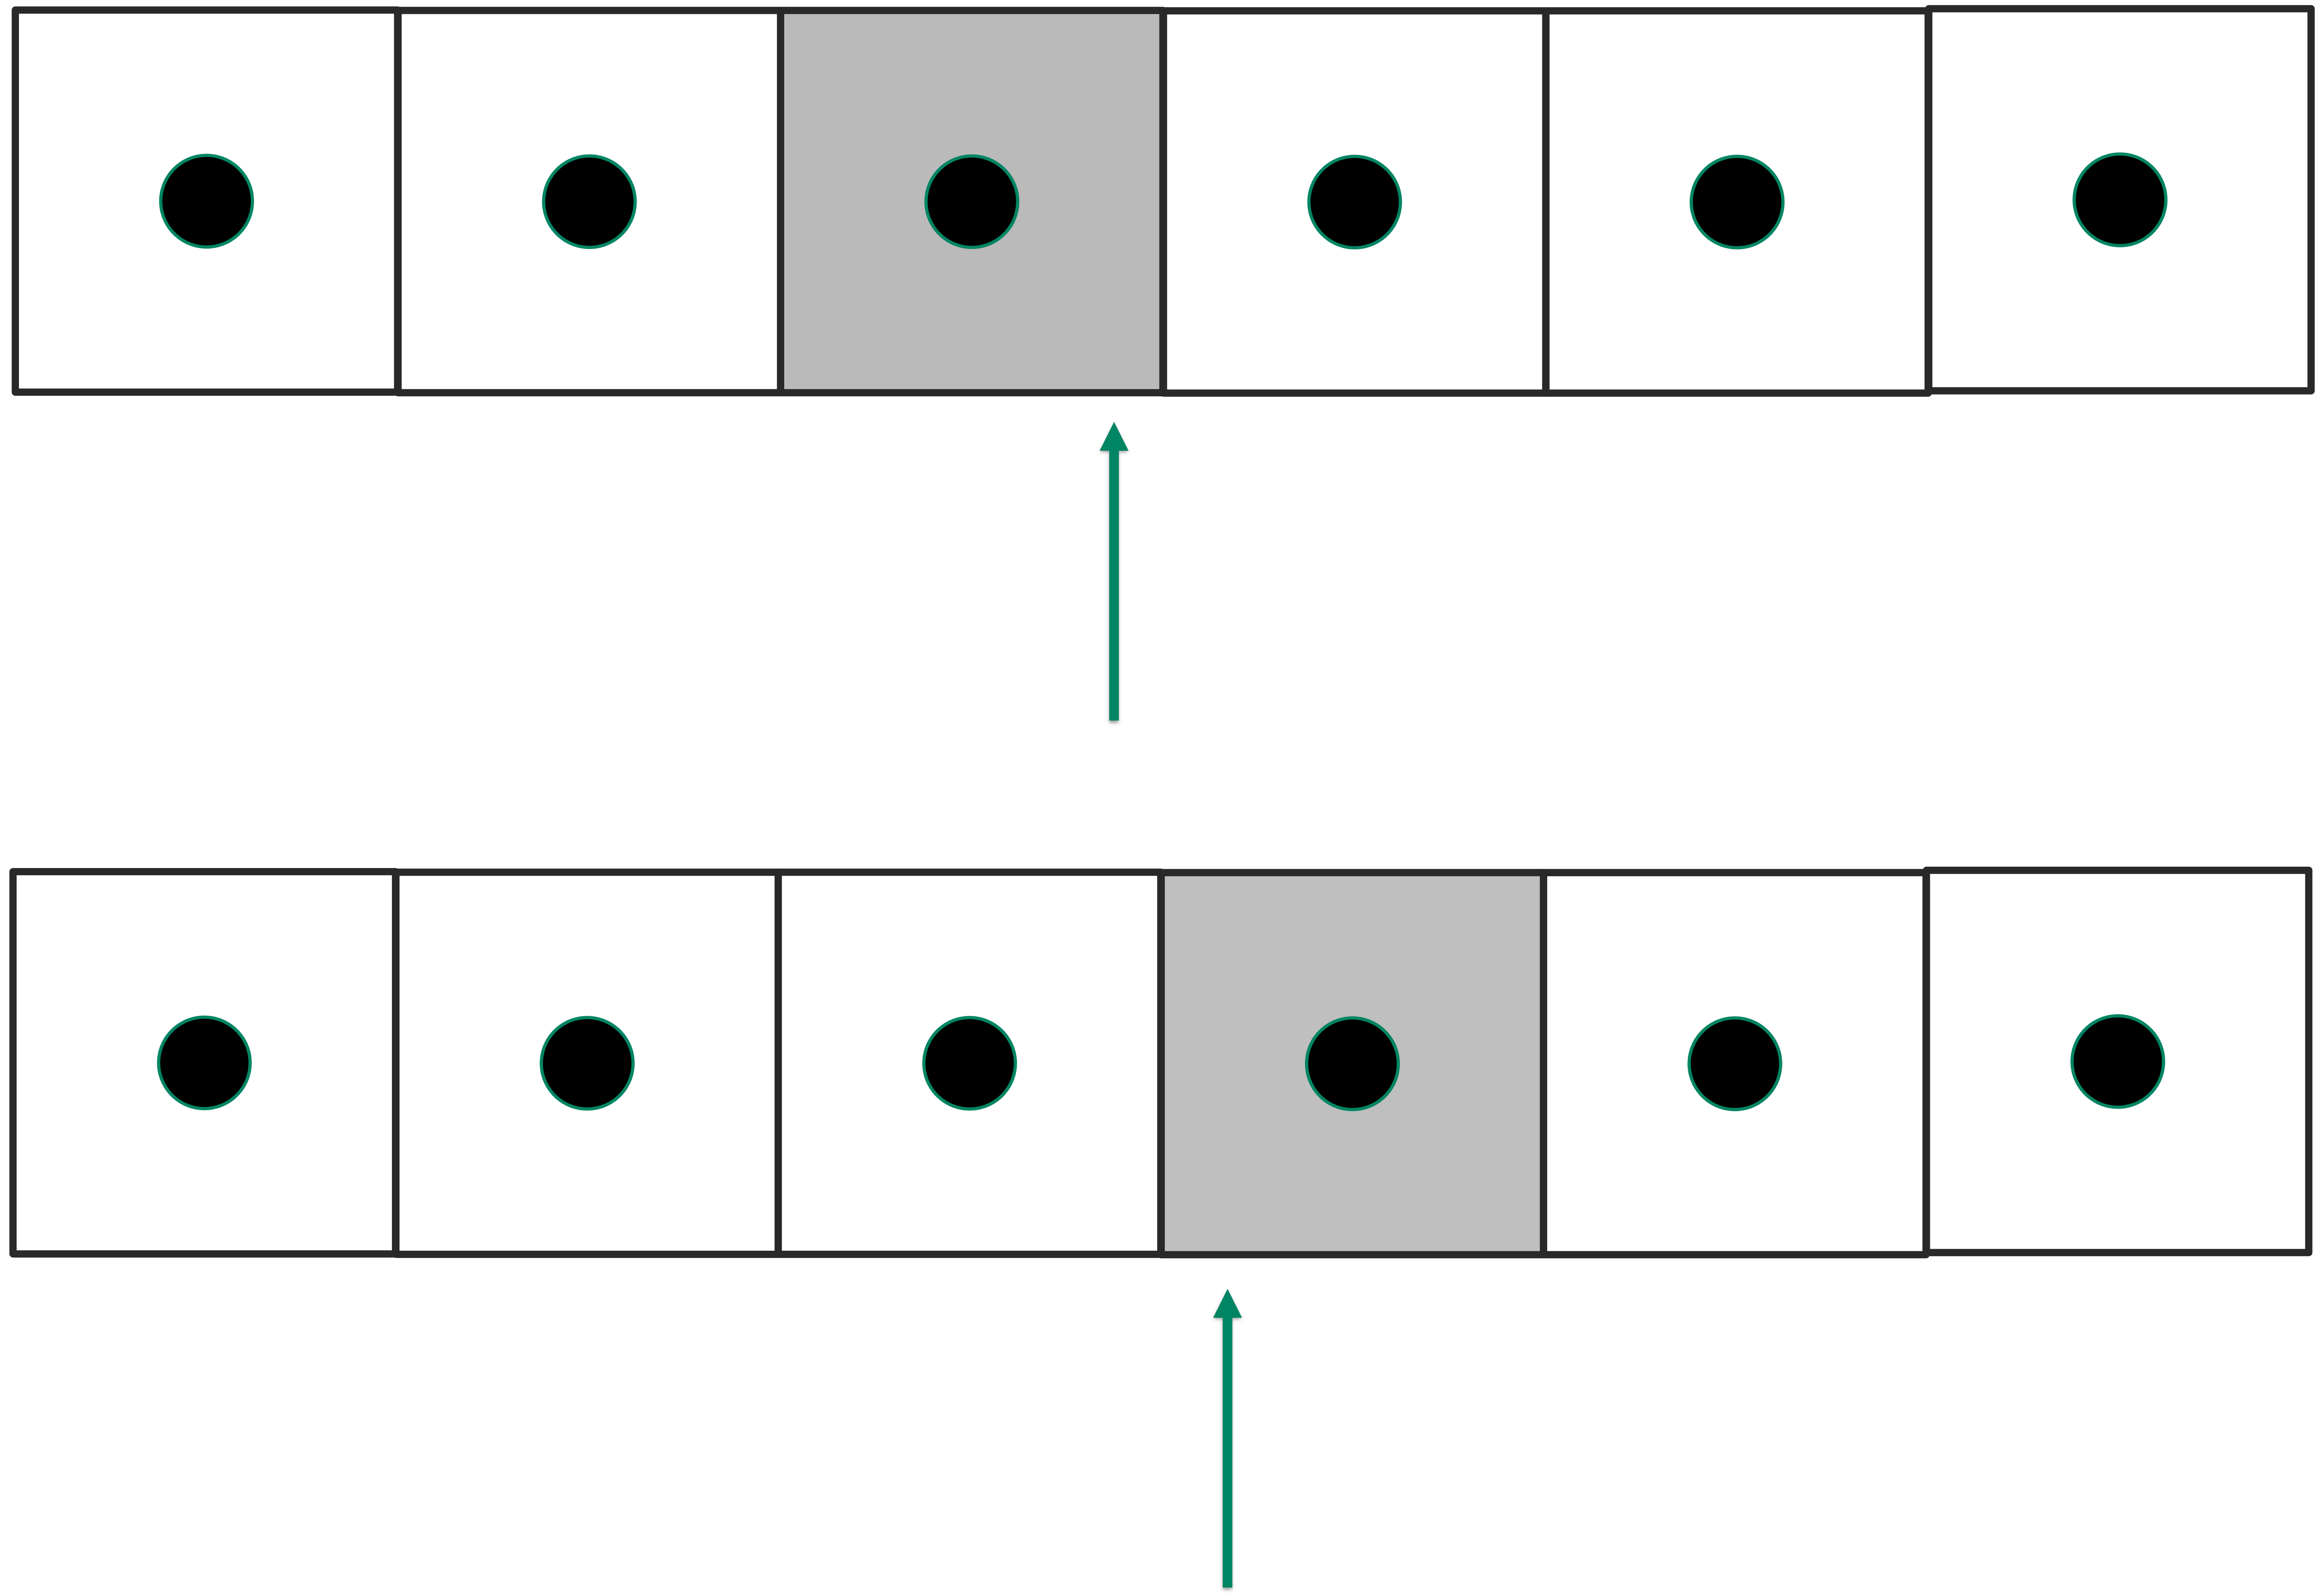

Magnification
Filter Importance Sampling through UV jittering

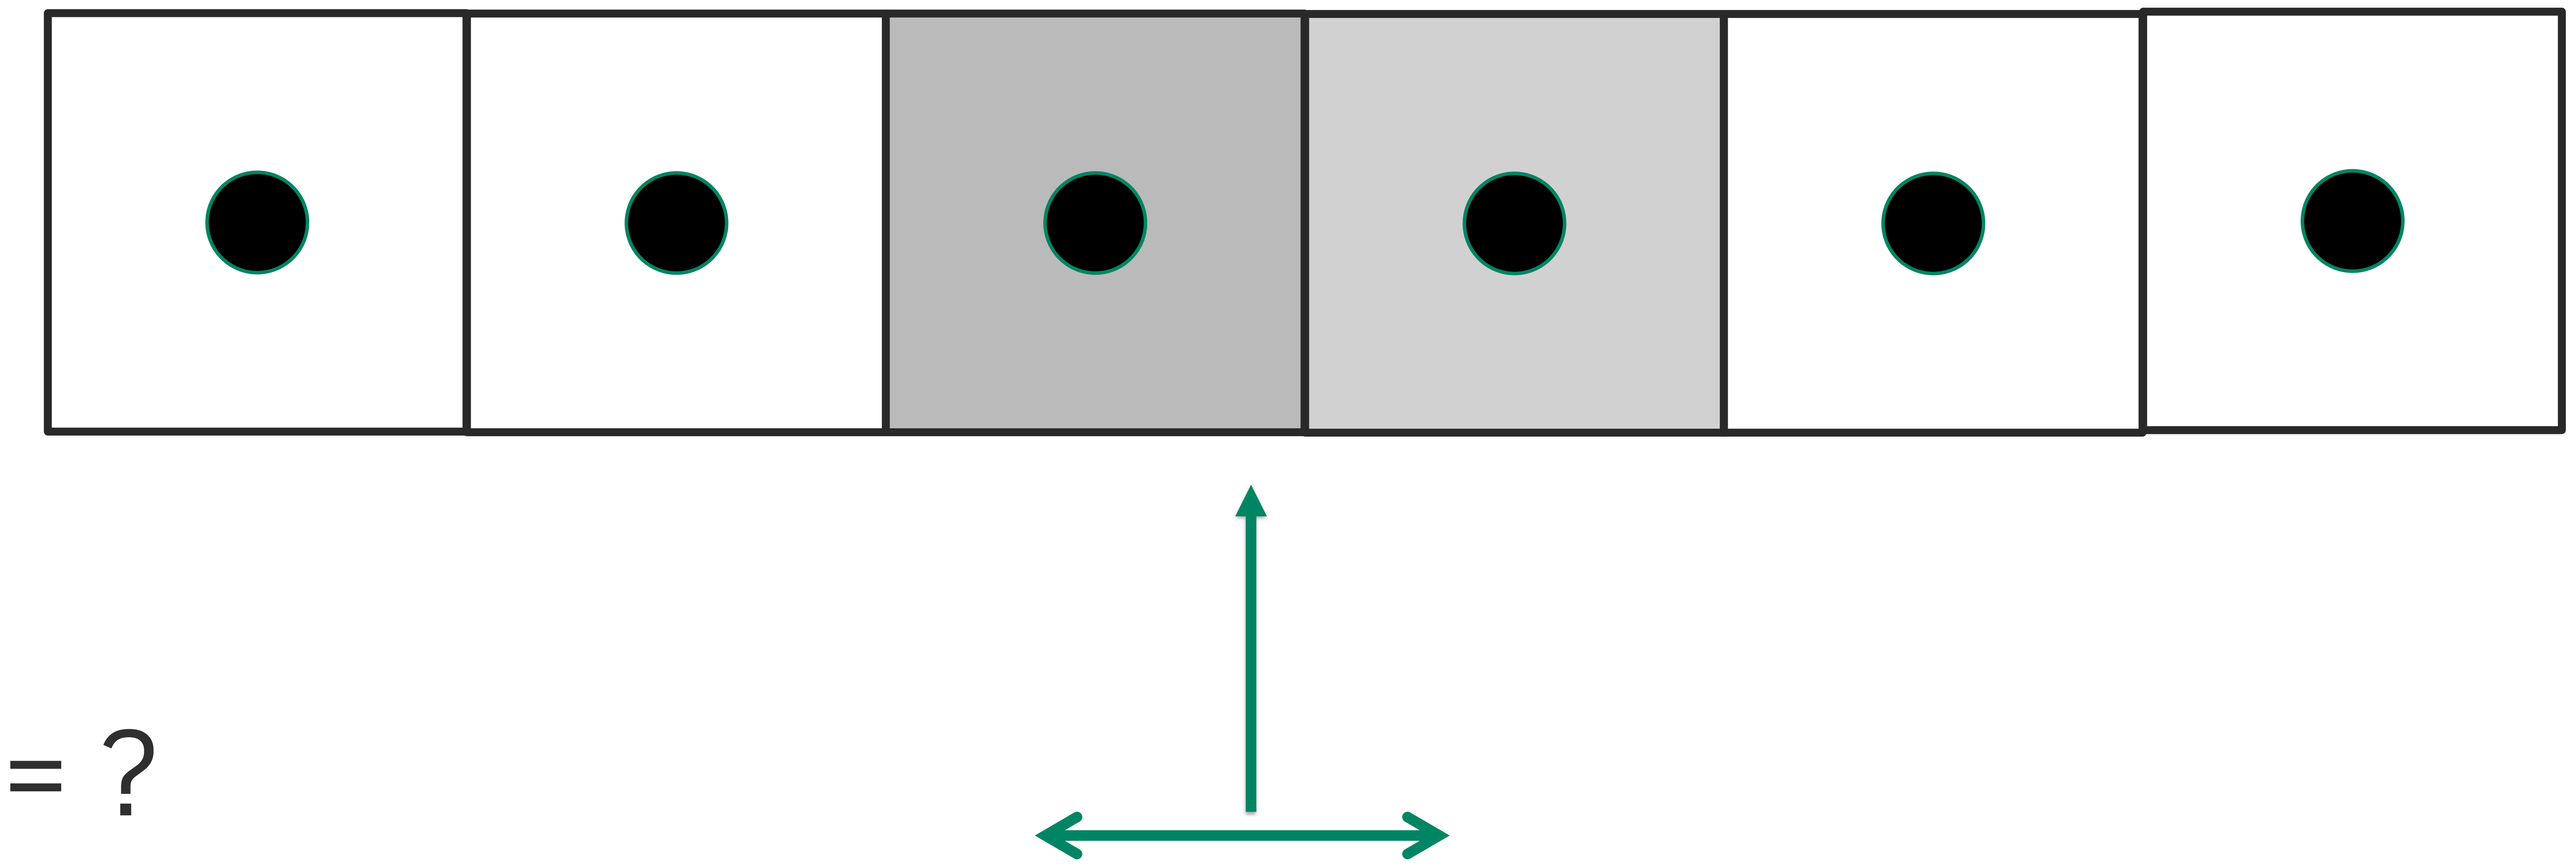

Uniform UV jitter + nearest neighbor = ?

Magnification
Filter Importance Sampling through UV jittering

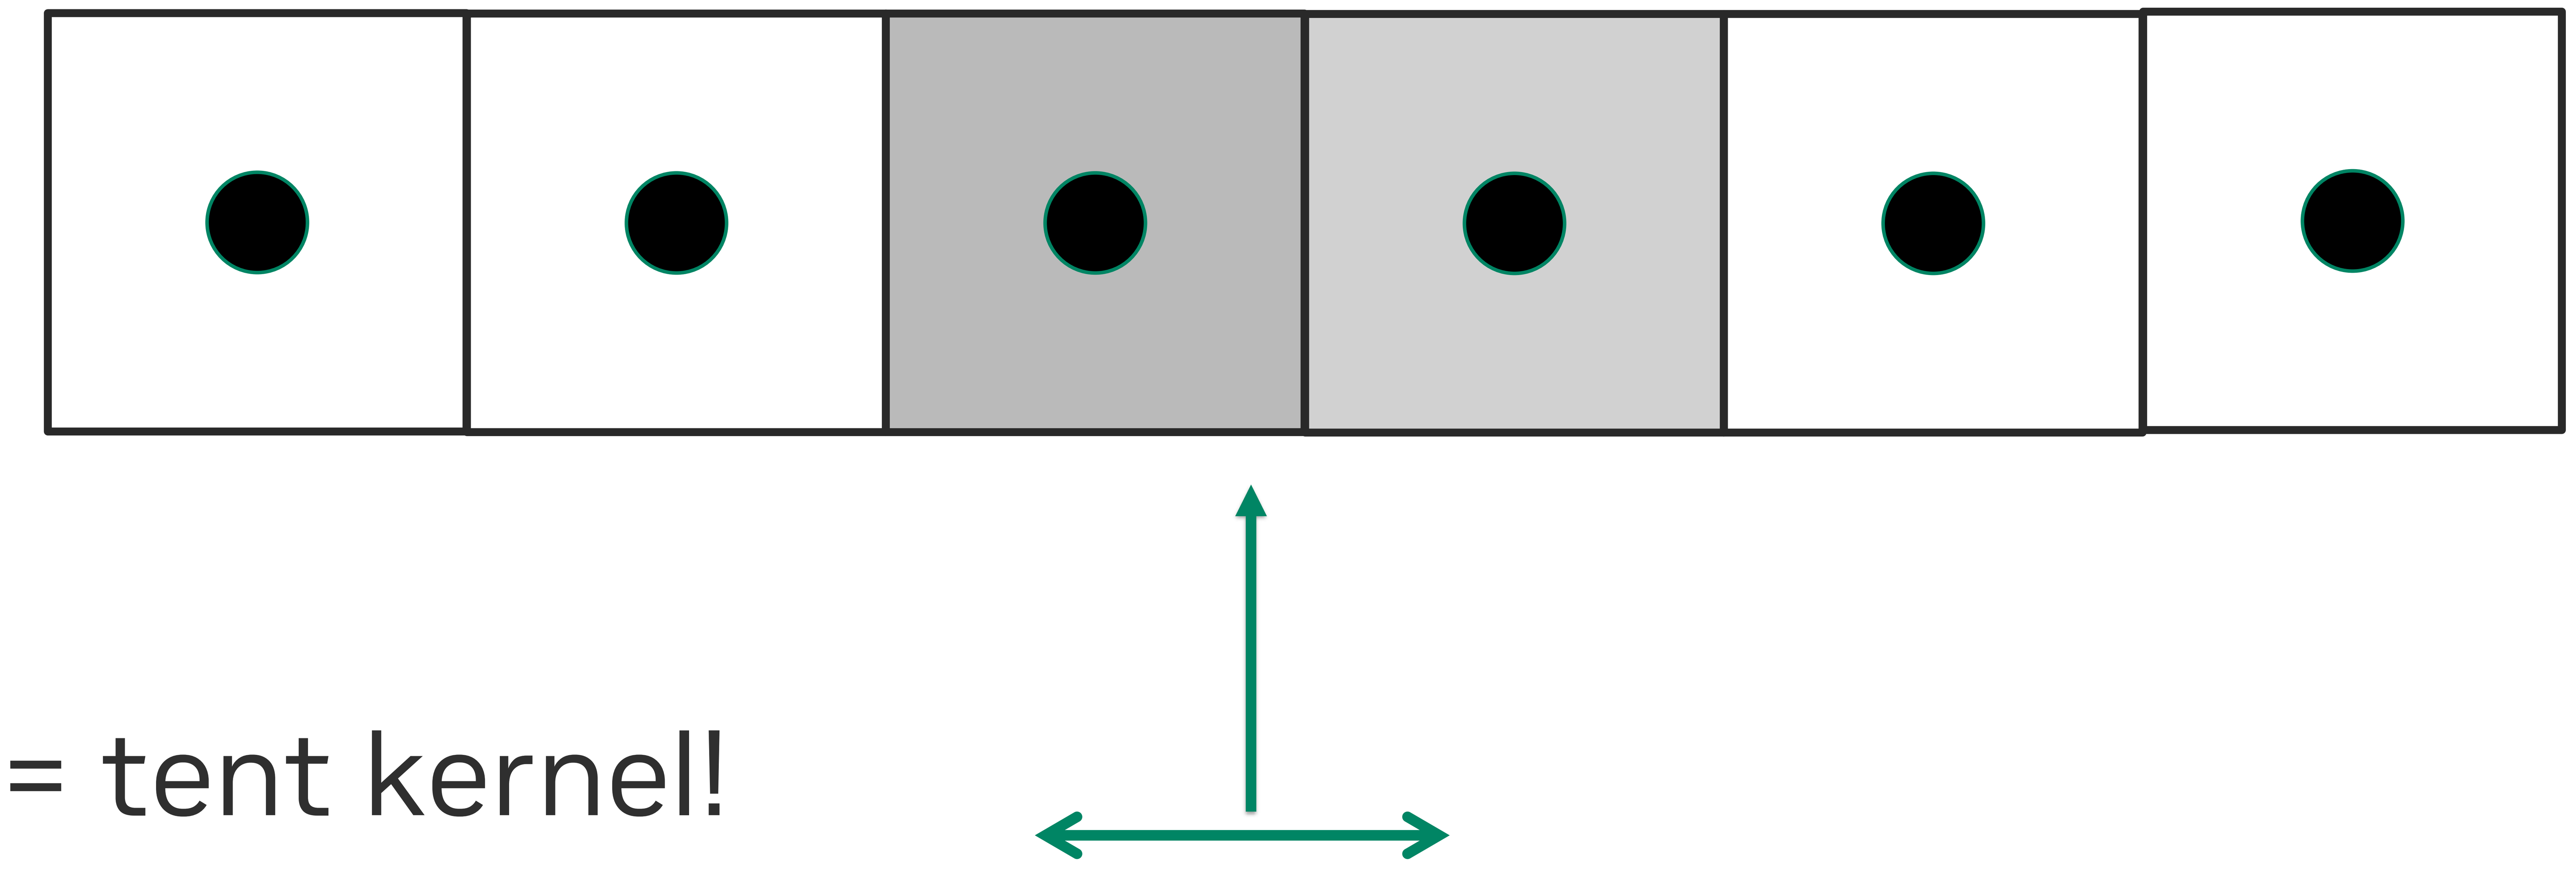

Uniform UV jitter + nearest neighbor = tent kernel!
The same as **linear interpolation**

Jitter PDF * Box Kernel Convolution

Magnification
Filter Importance Sampling through UV jittering

Linear/tent UV jitter + nearest neighbor box = quadratic B-Spline
Quadratic UV jitter + nearest neighbor box = **cubic B-Spline**

Magnification
Filter Importance Sampling through UV jittering

- For B-Spline filters, this additional box is desirable!
- Can sample other, including infinite spatial support filters
- Jitter UVs according to PDF deconvolved with a box

Magnification
Filter Importance Sampling through UV jittering

- For B-Spline filters, this additional box is desirable!
- Can sample other, including infinite spatial support filters
- Jitter UVs according to PDF deconvolved with a box
- For many other filters -> can still be advantageous (prevent Gaussian under-sampling)

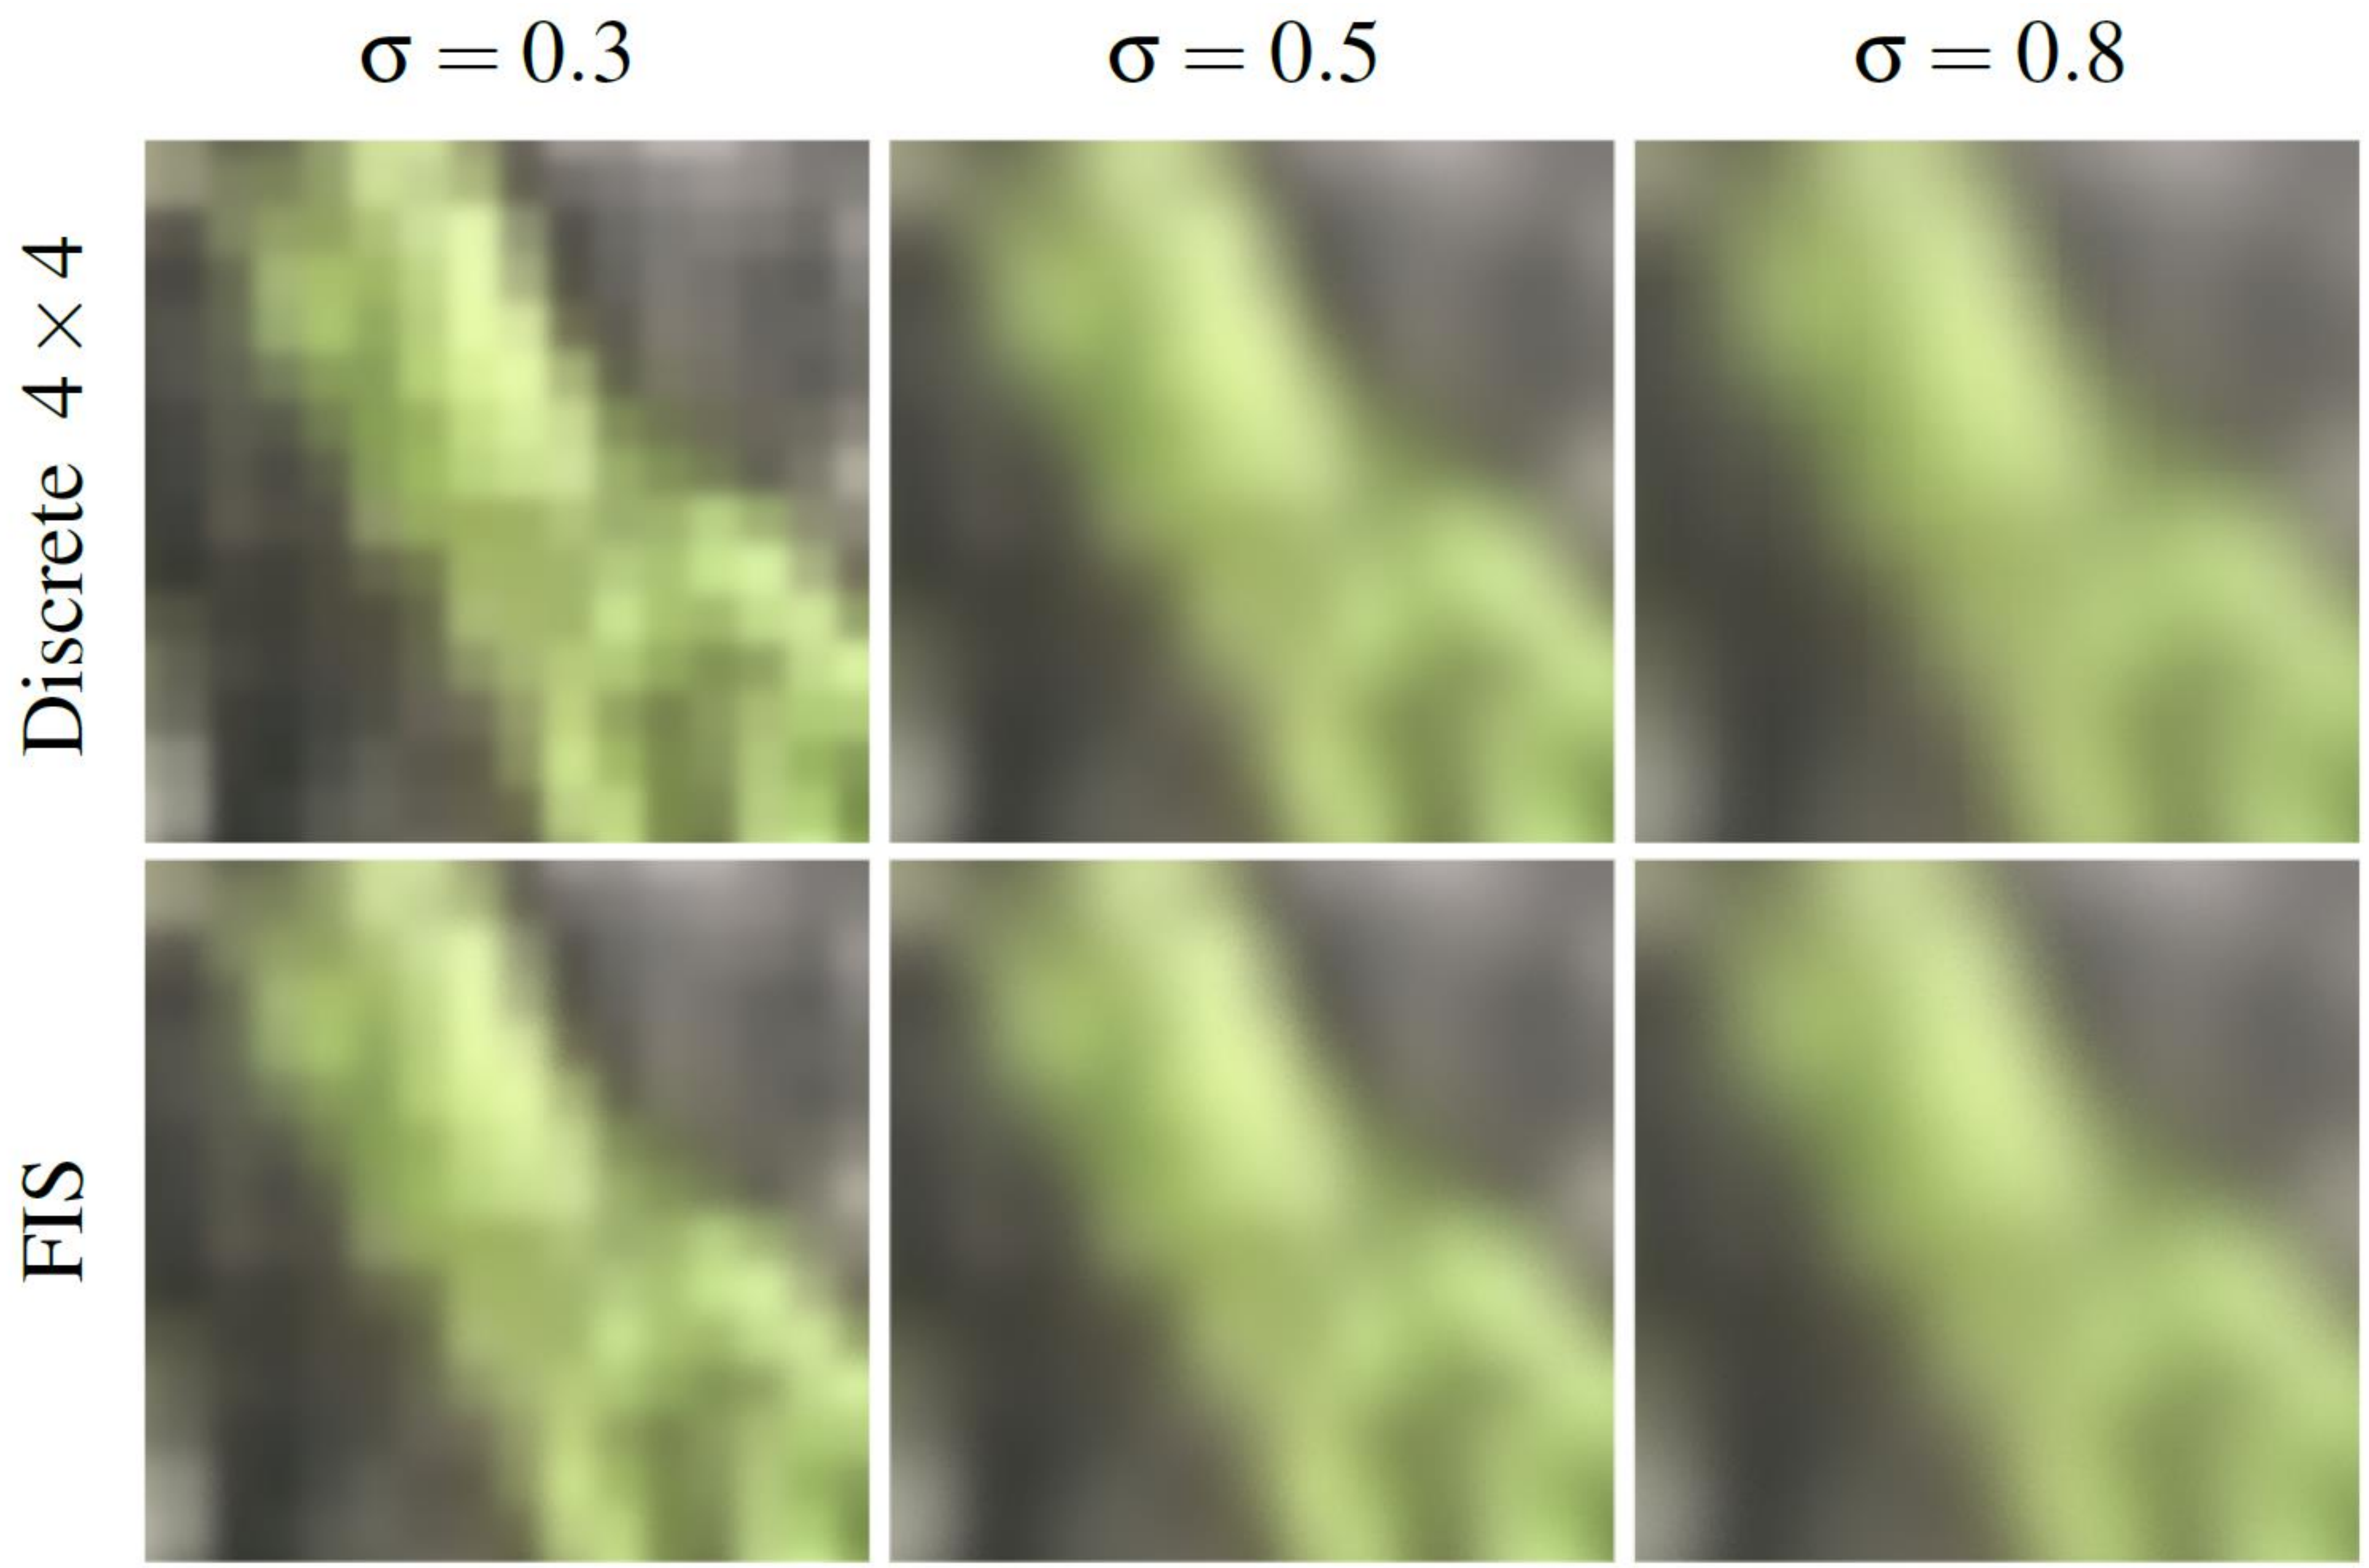

Magnification
Filter Importance Sampling through UV jittering

- For B-Spline filters, this additional box is desir
- Can sample other, including infinite spatial su
- Jitter UVs according to PDF deconvolved with
- For many other filters -> can still be advantag

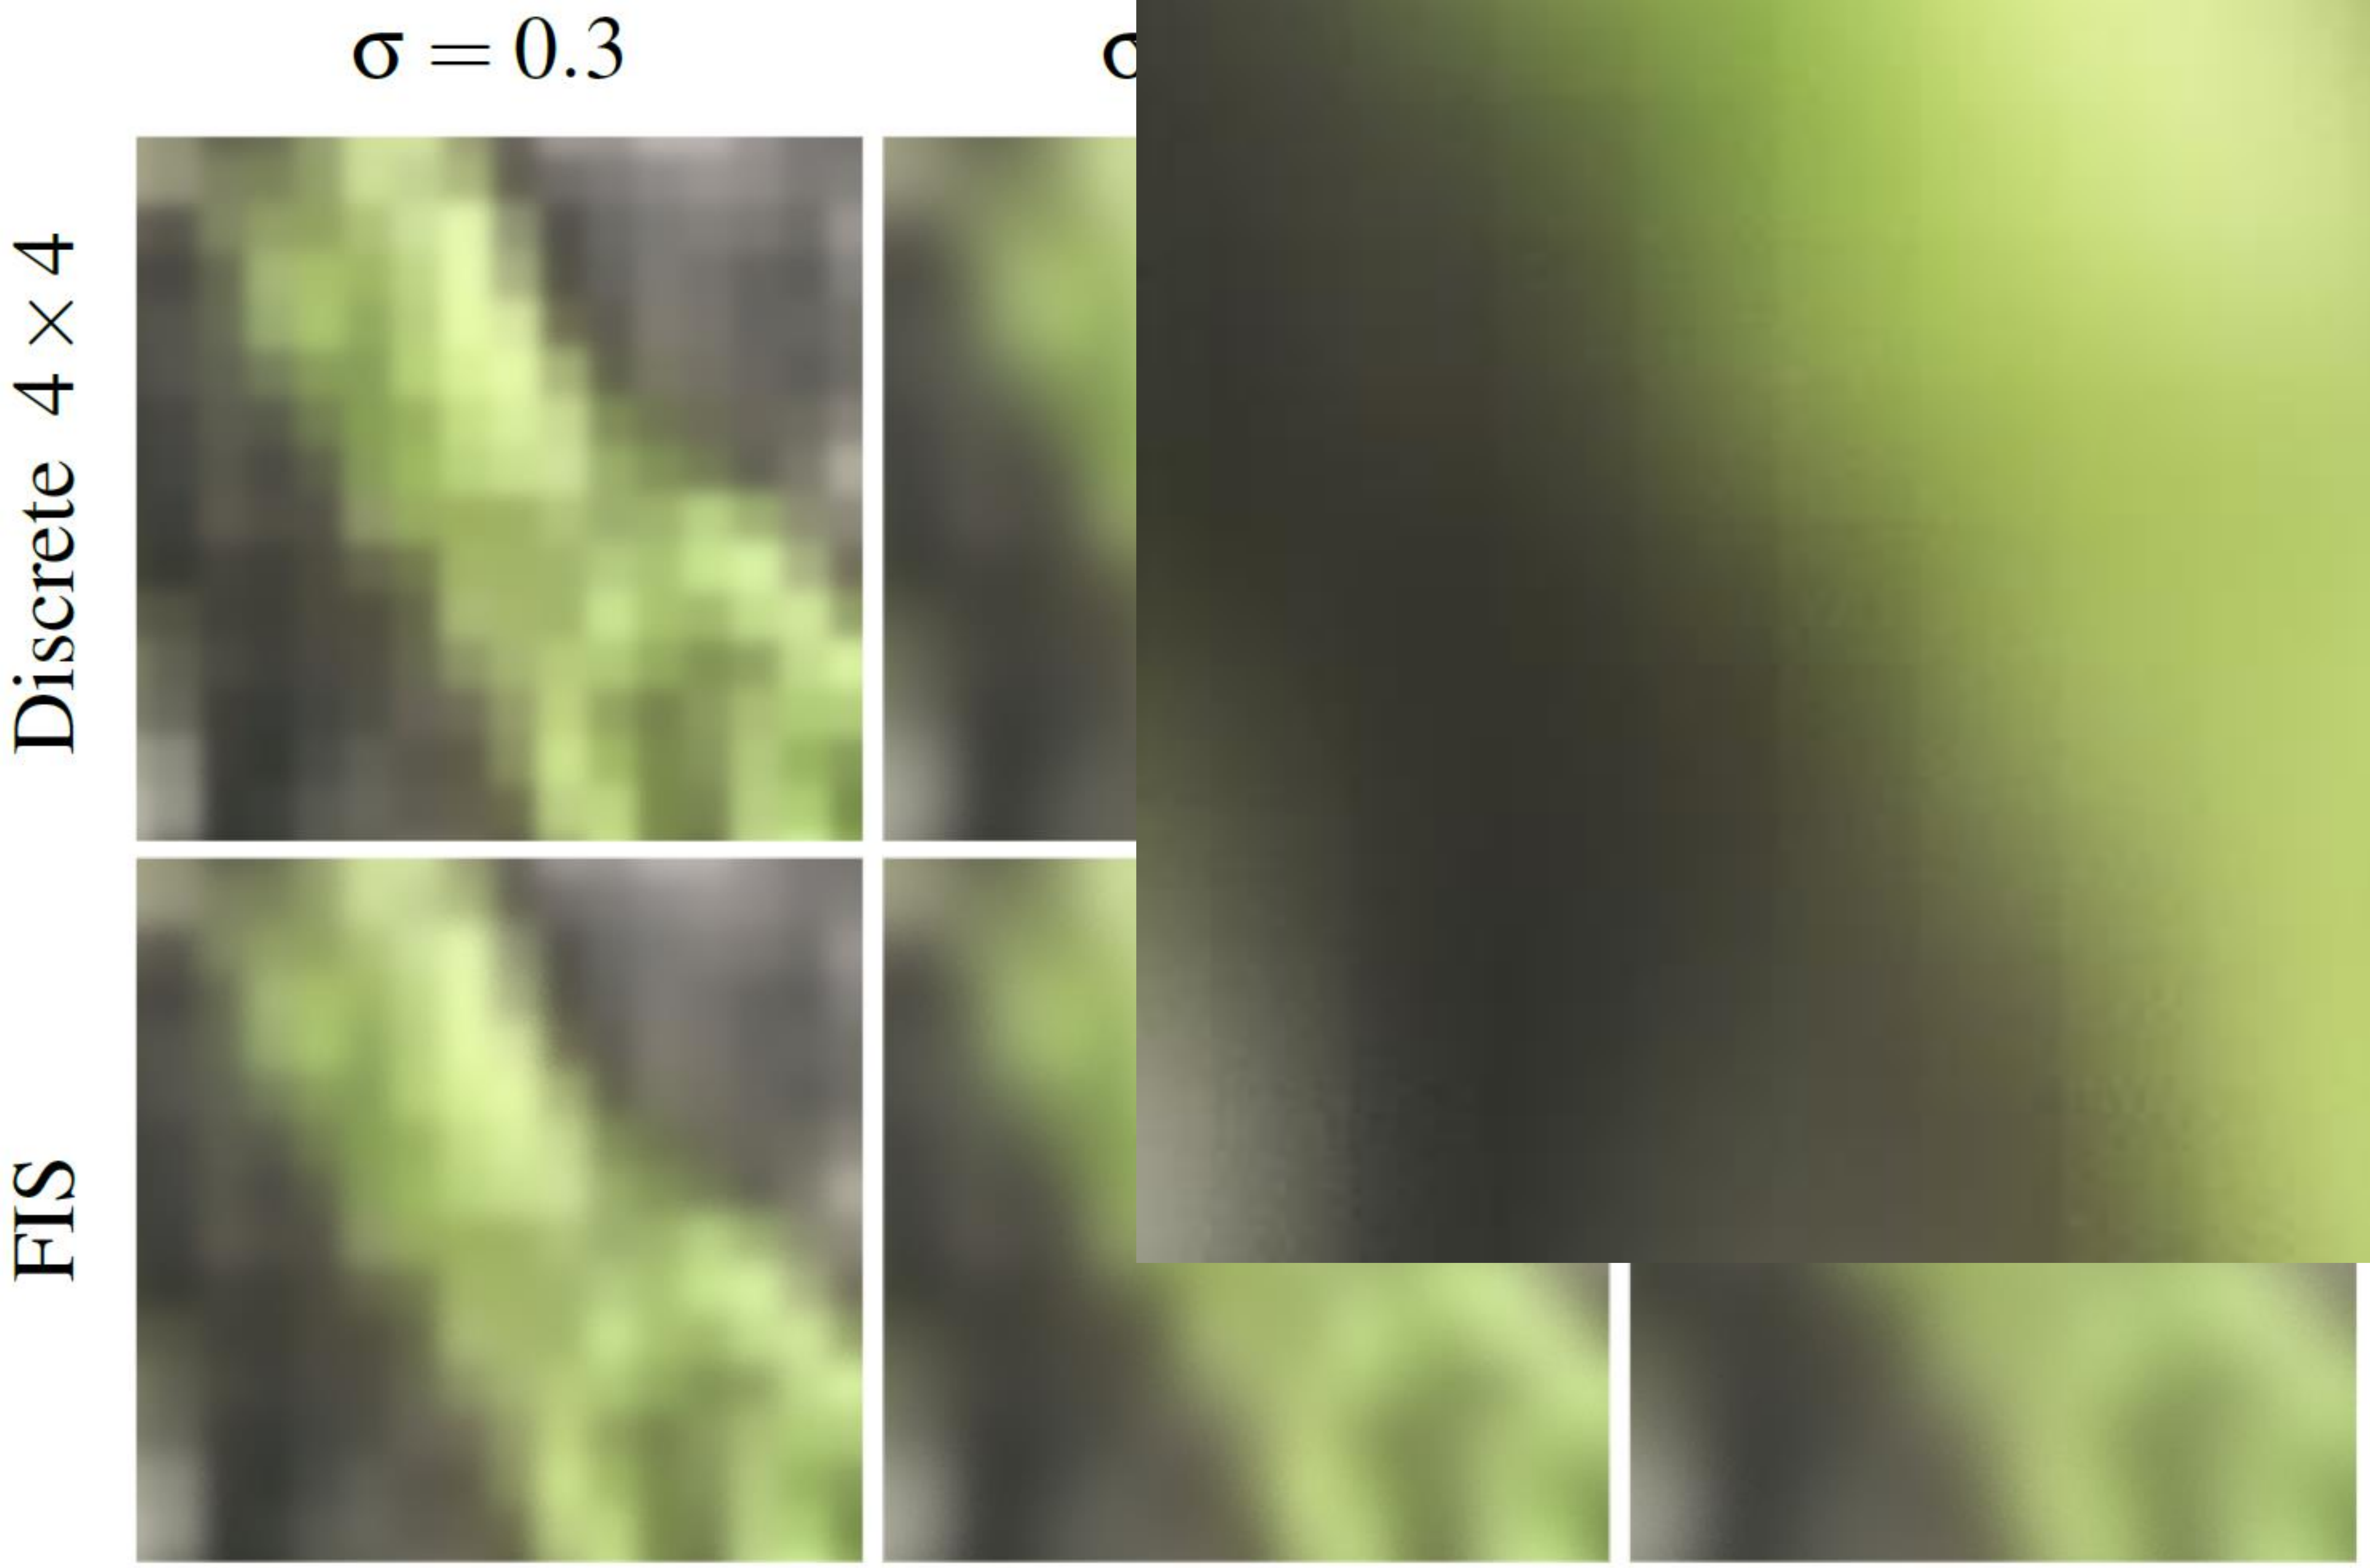

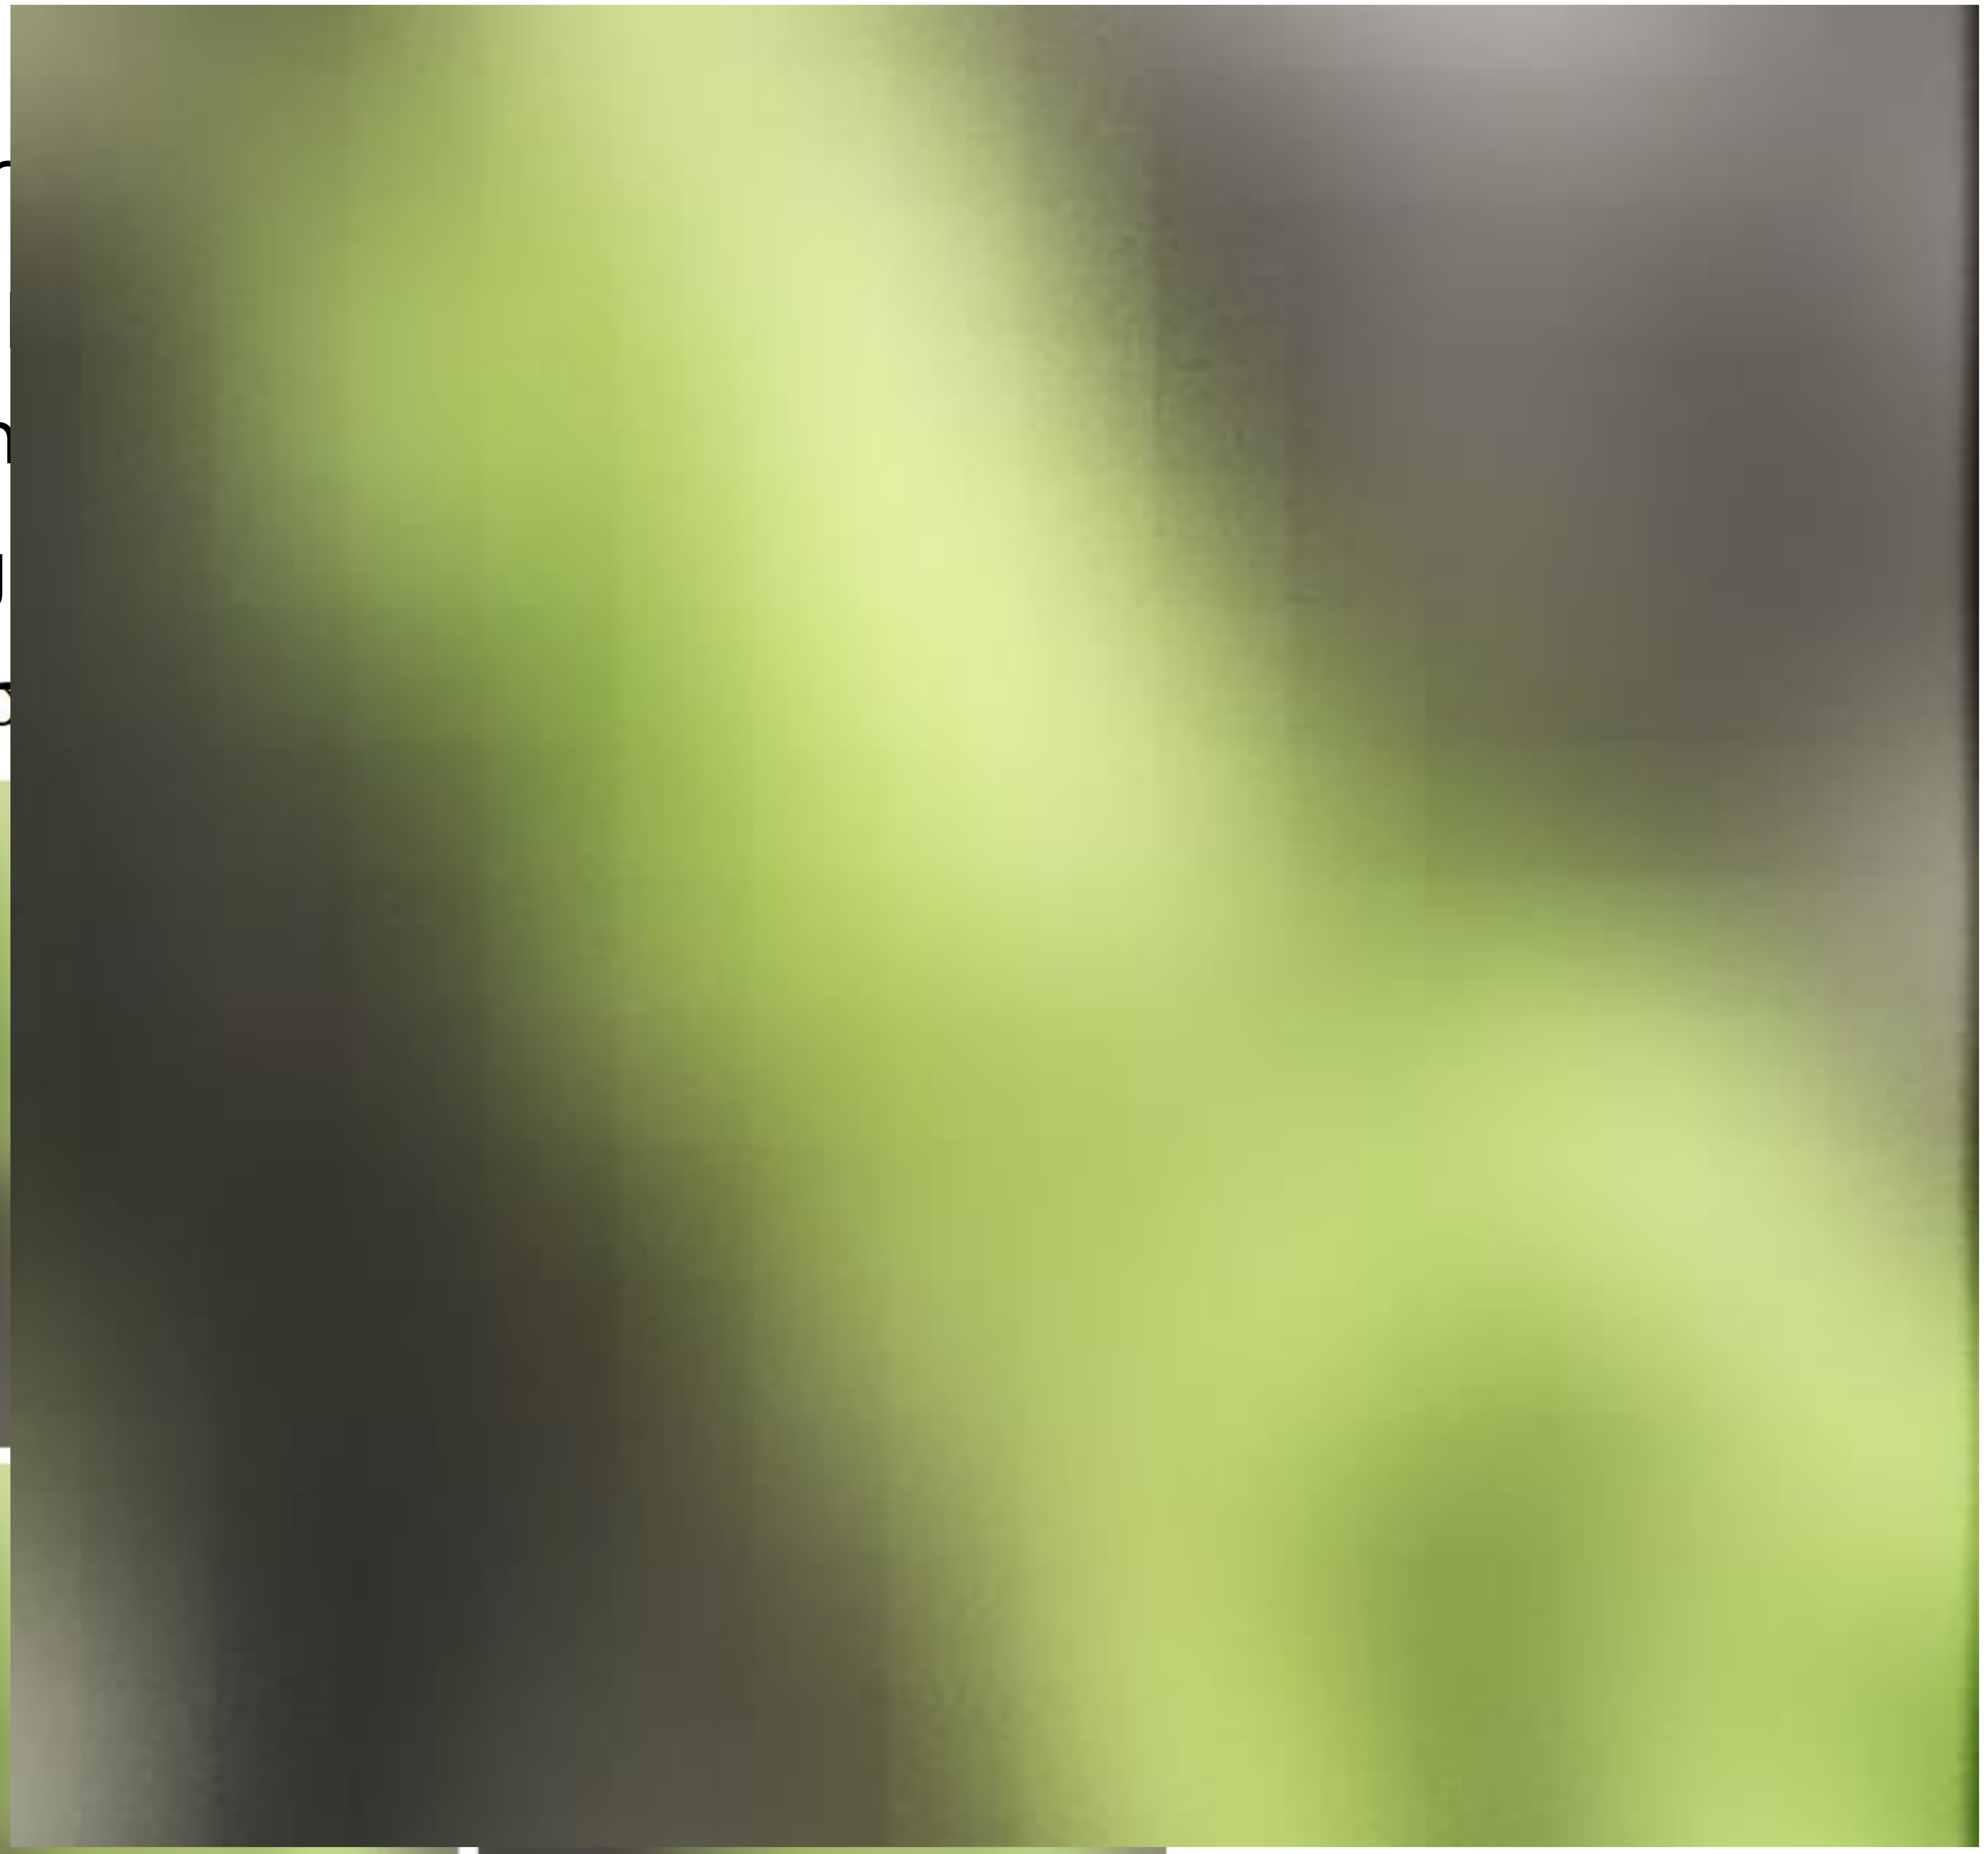

Stochastic Filtering families compared

- Main difference: **discrete vs continuous domain**

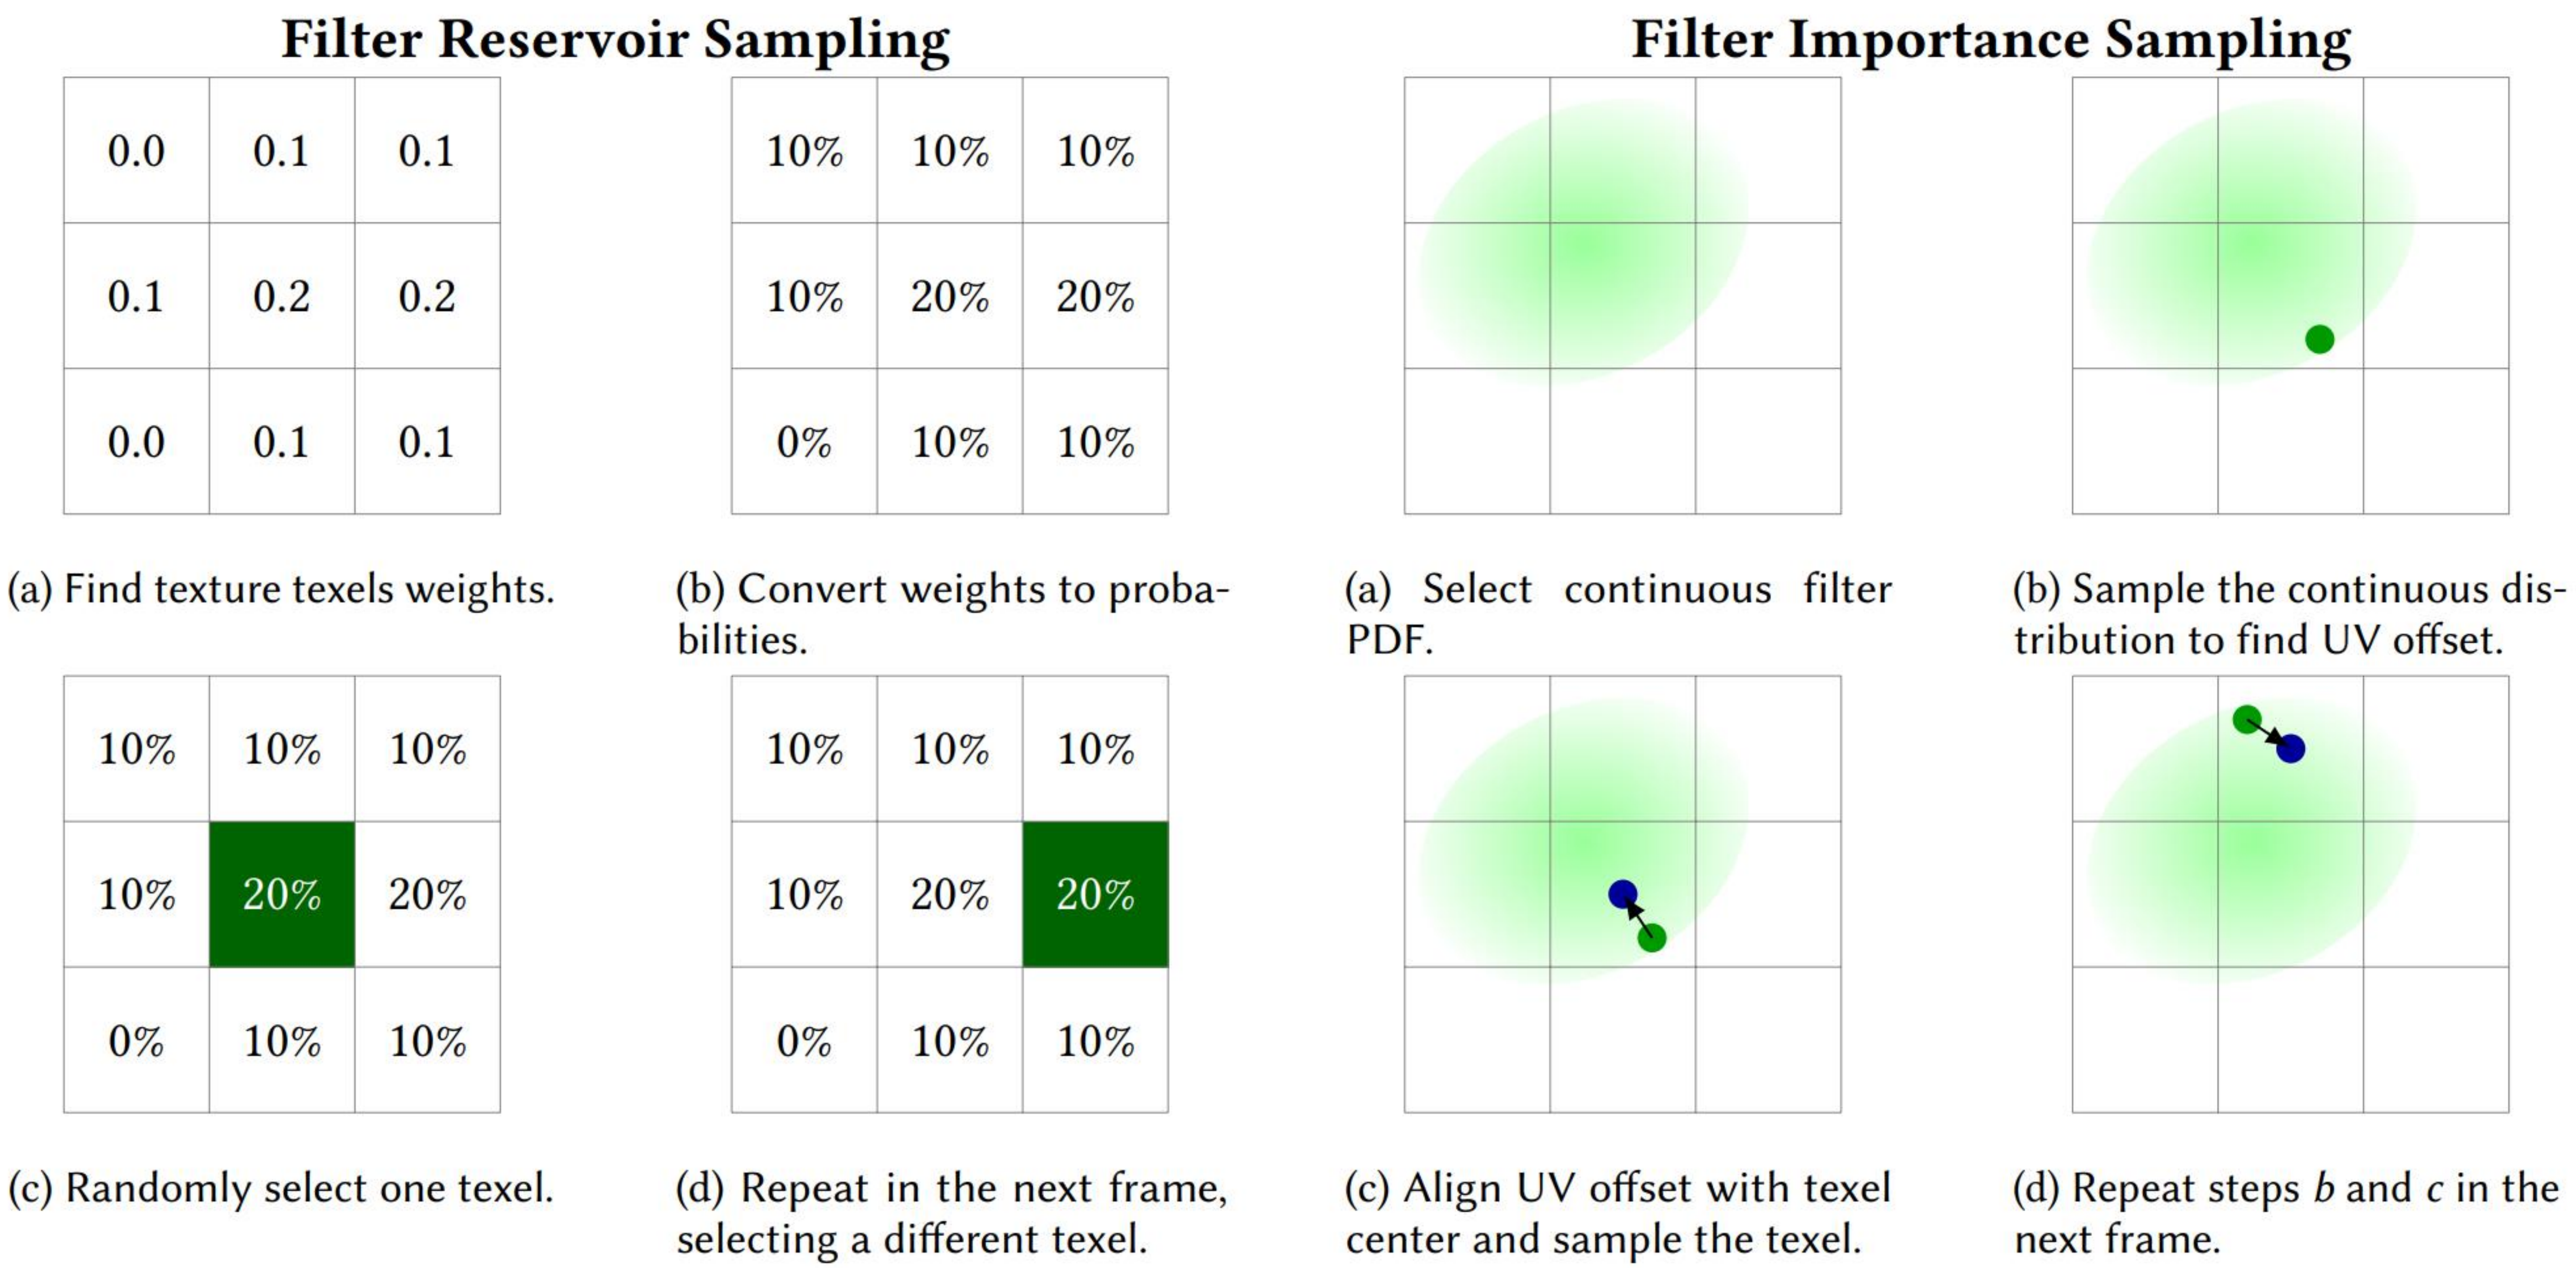

Stochastic Filtering families compared

- Main difference: **discrete vs continuous domain**
- In many cases, FRS is the only option (arbitrary discrete kernels, positivization)
- Otherwise, we recommend FIS – simpler implementation, see provided source code

Filter Reservoir Sampling | **Filter Importance Sampling**

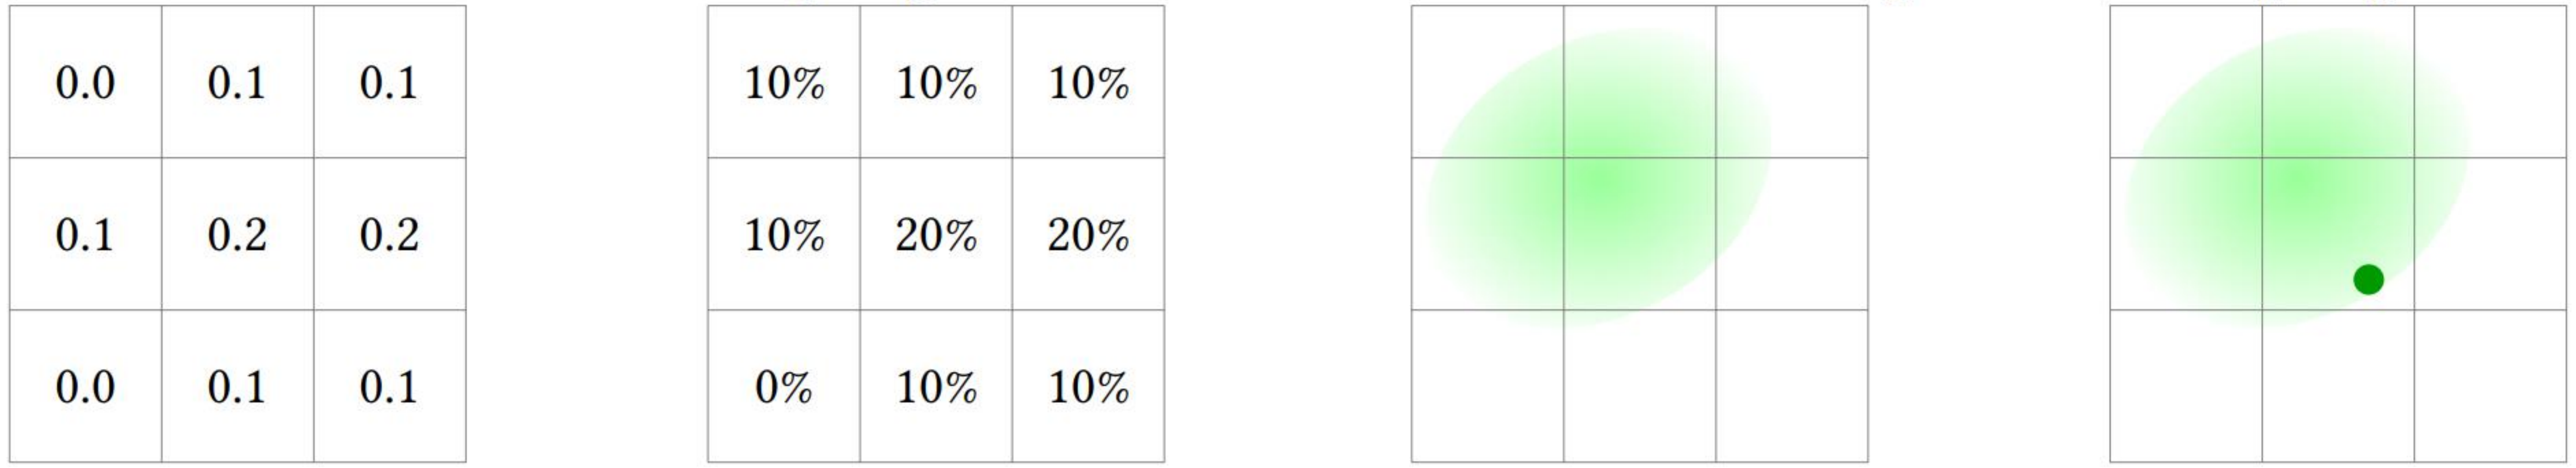

(a) Find texture texels weights. (b) Convert weights to probabilities. (a) Select continuous filter PDF. (b) Sample the continuous distribution to find UV offset.

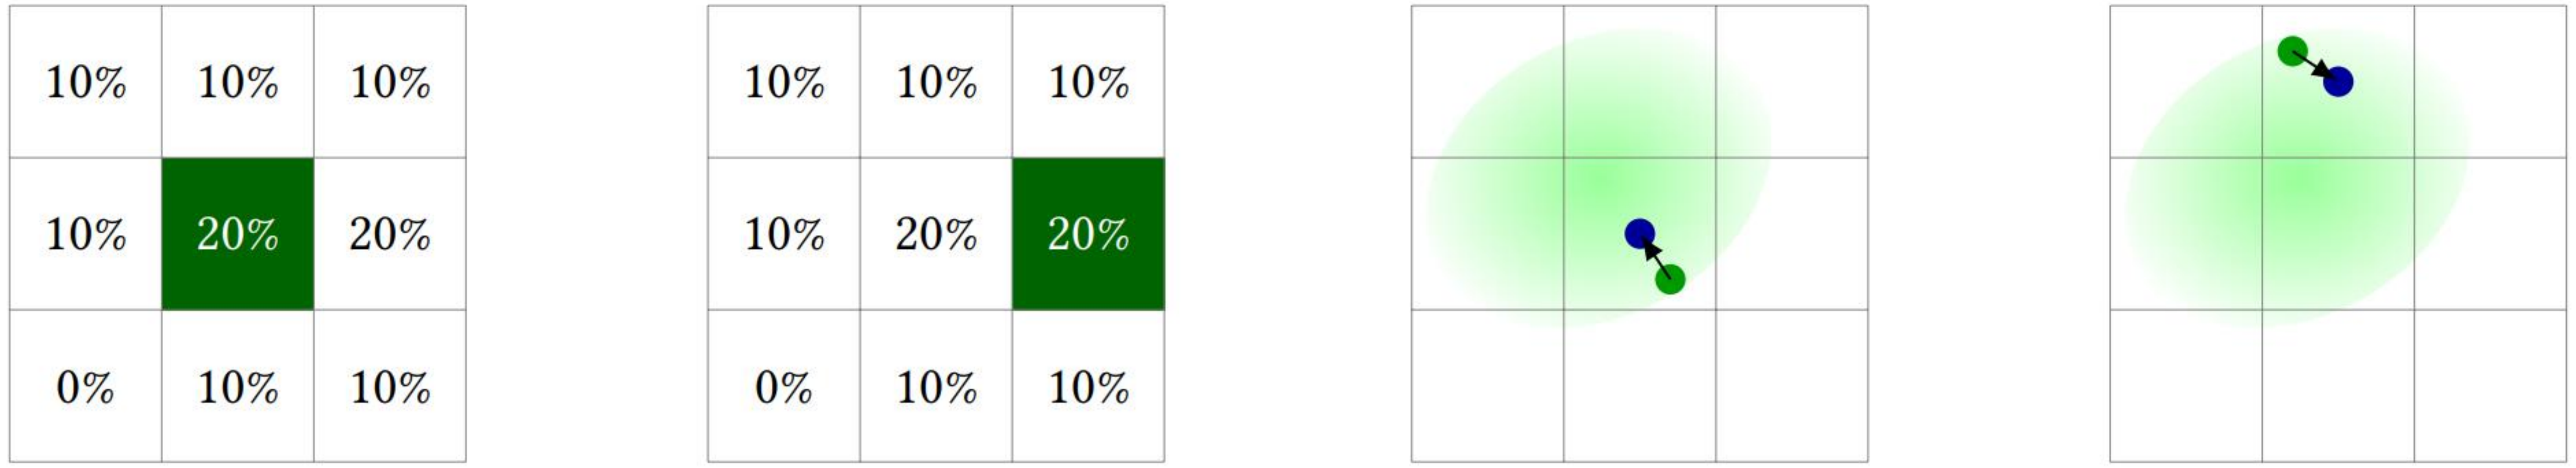

(c) Randomly select one texel. (d) Repeat in the next frame, selecting a different texel. (c) Align UV offset with texel center and sample the texel. (d) Repeat steps *b* and *c* in the next frame.

Minification
Stochastic Filtering After Shading

- Anisotropic filtering or elliptically weighted average
- Many pixels, non-uniform mapping for jittering
- There's a simpler, **already-used** method!

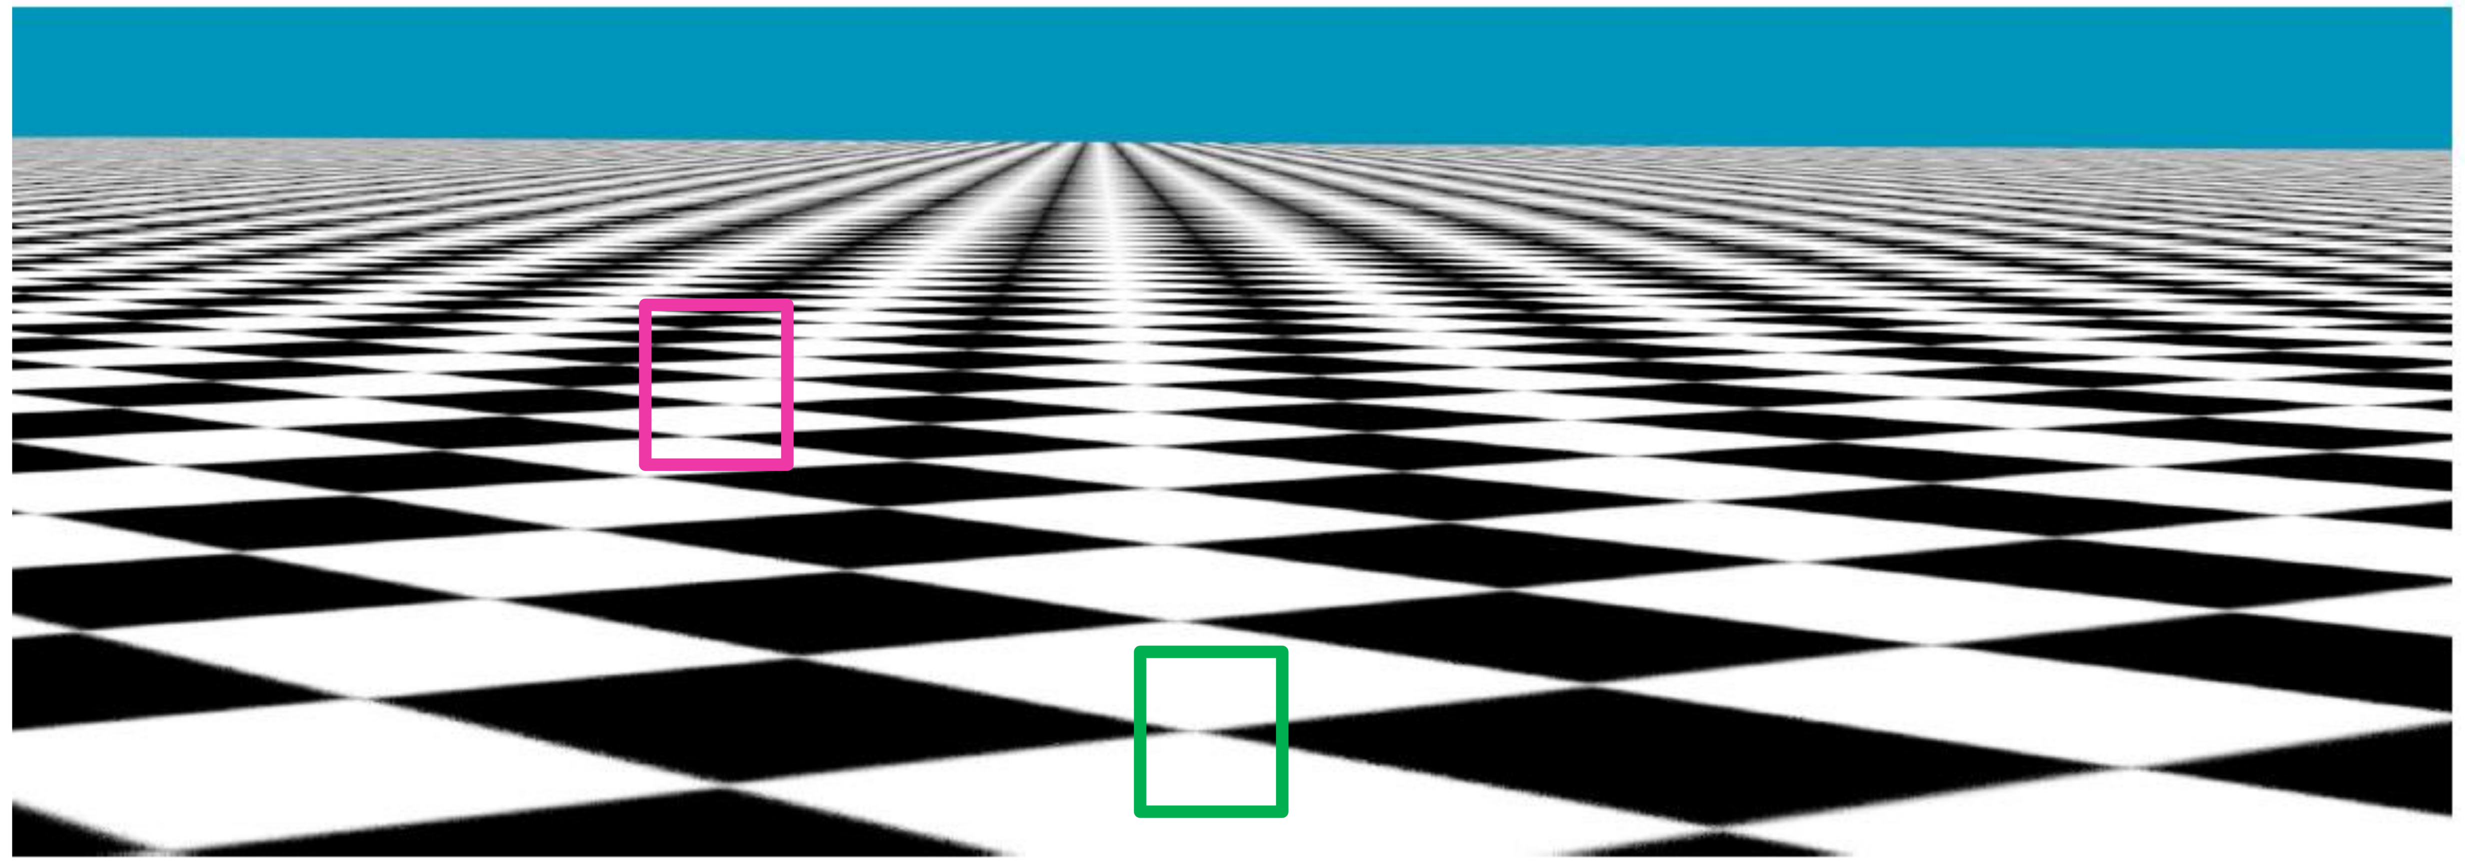

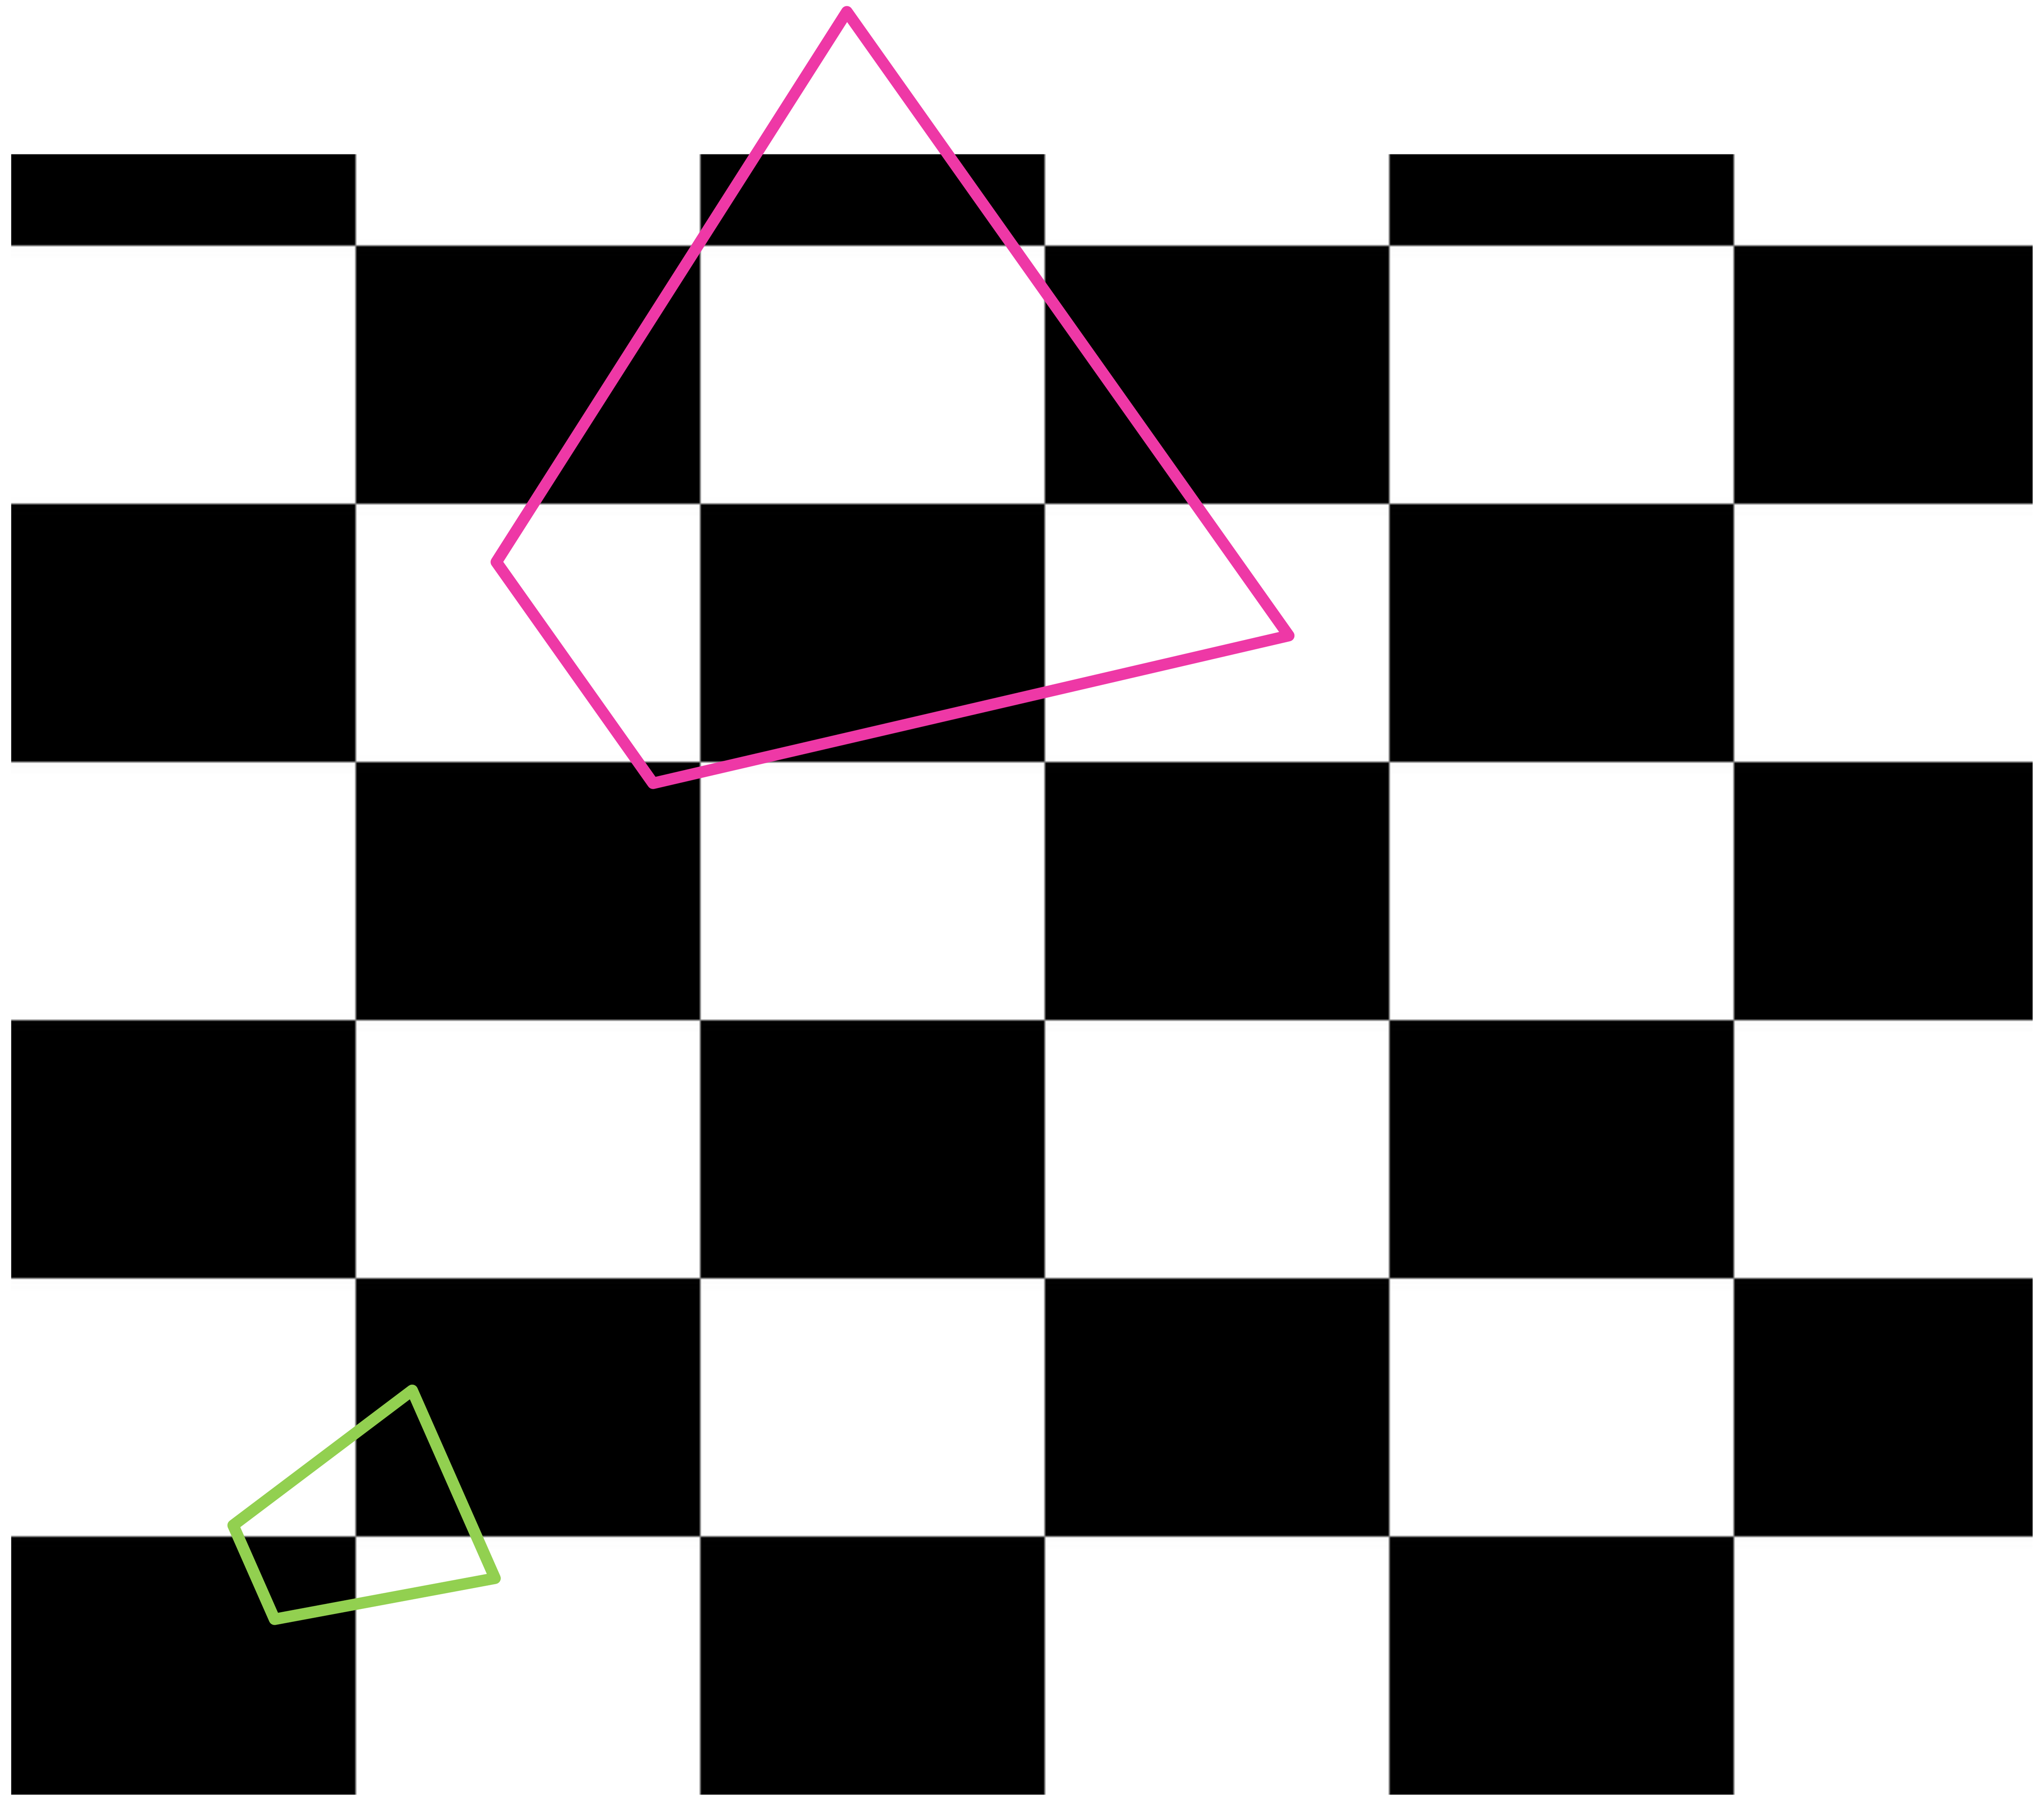

Minification
Stochastic Filtering After Shading

- Common practice – jitter the projection matrix for anti-aliasing reconstruction filter
- Used offline (e.g., MoonRay) and real-time (TAA, DLSS)

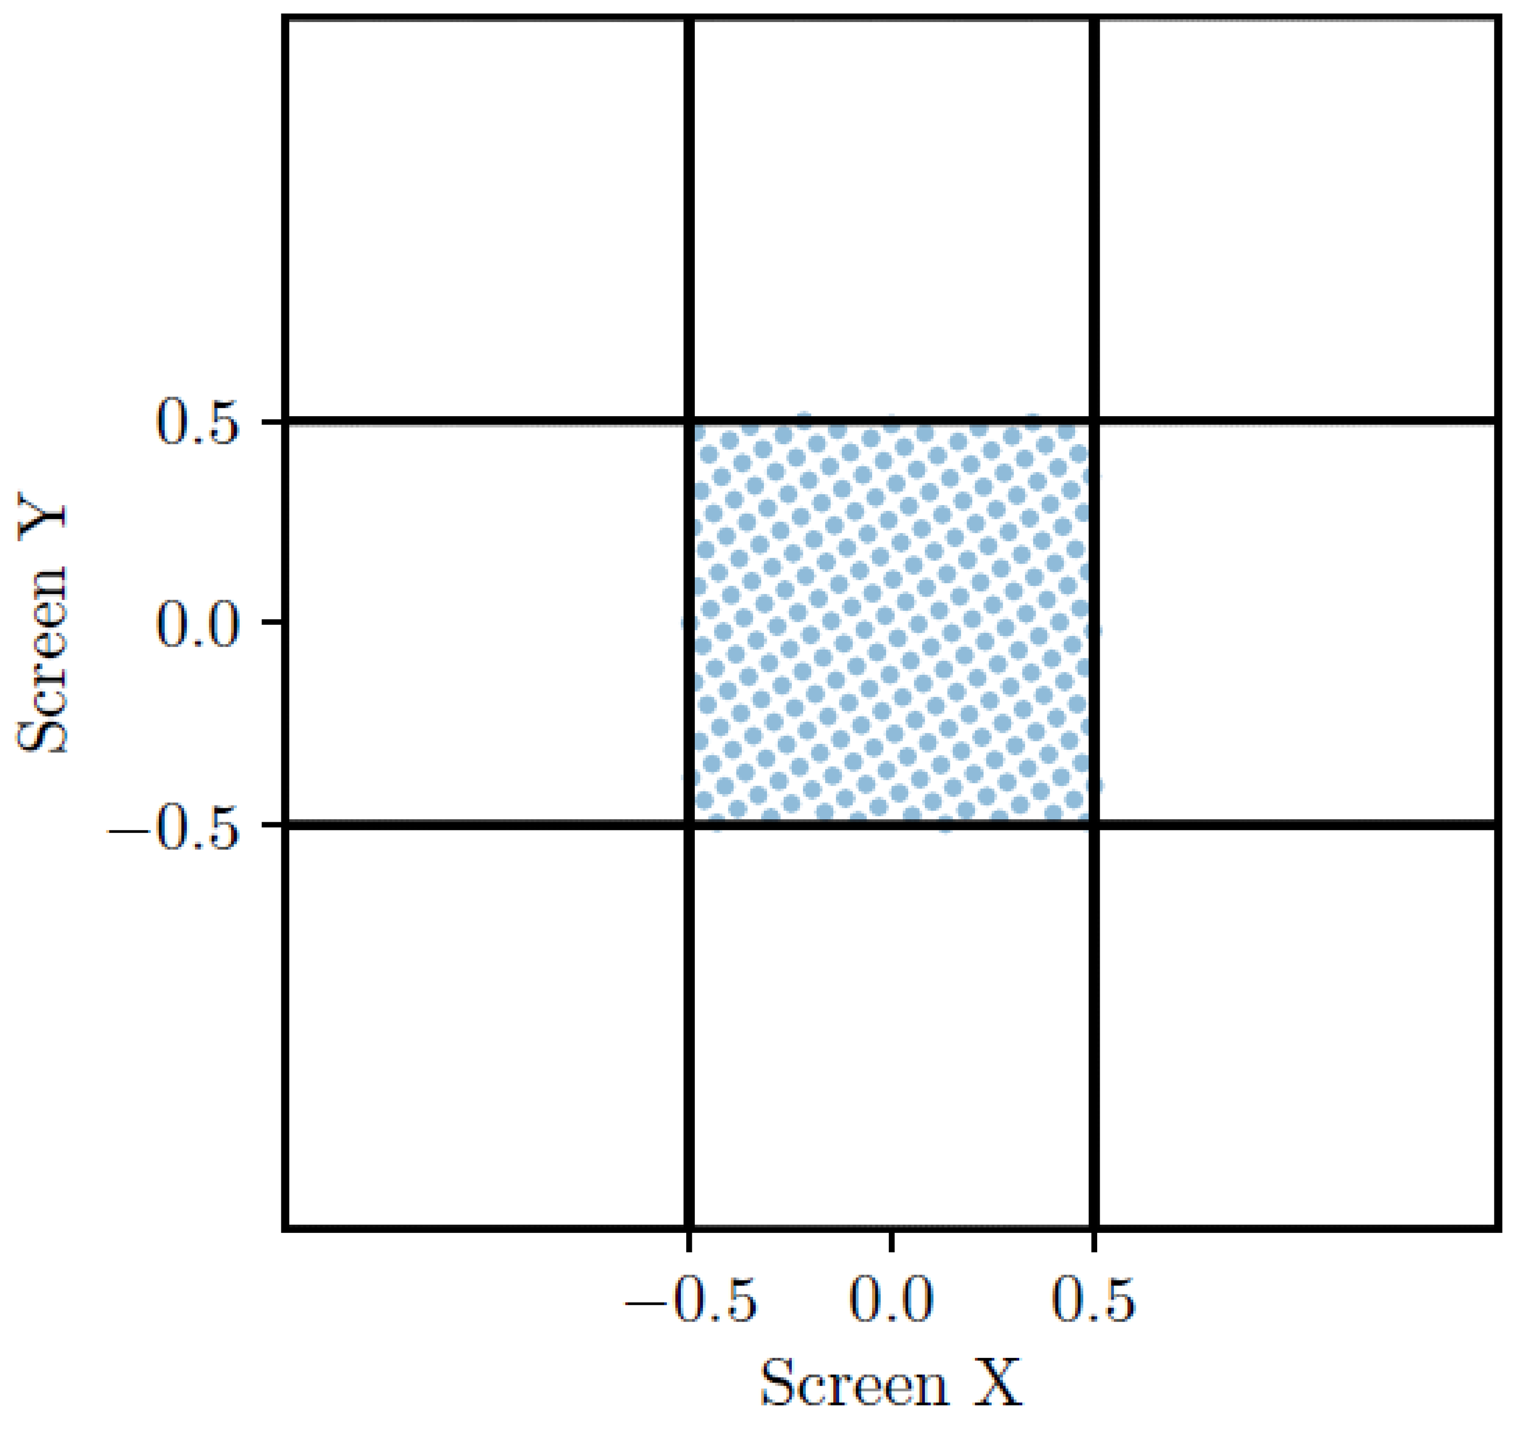

Minification
Stochastic Filtering After Shading

- Common practice – jitter the projection matrix for anti-aliasing reconstruction filter
- Used offline (e.g., MoonRay) and real-time (TAA, DLSS)
- Projects to trapezoid, minification supersampling -> **filtering after shading**!

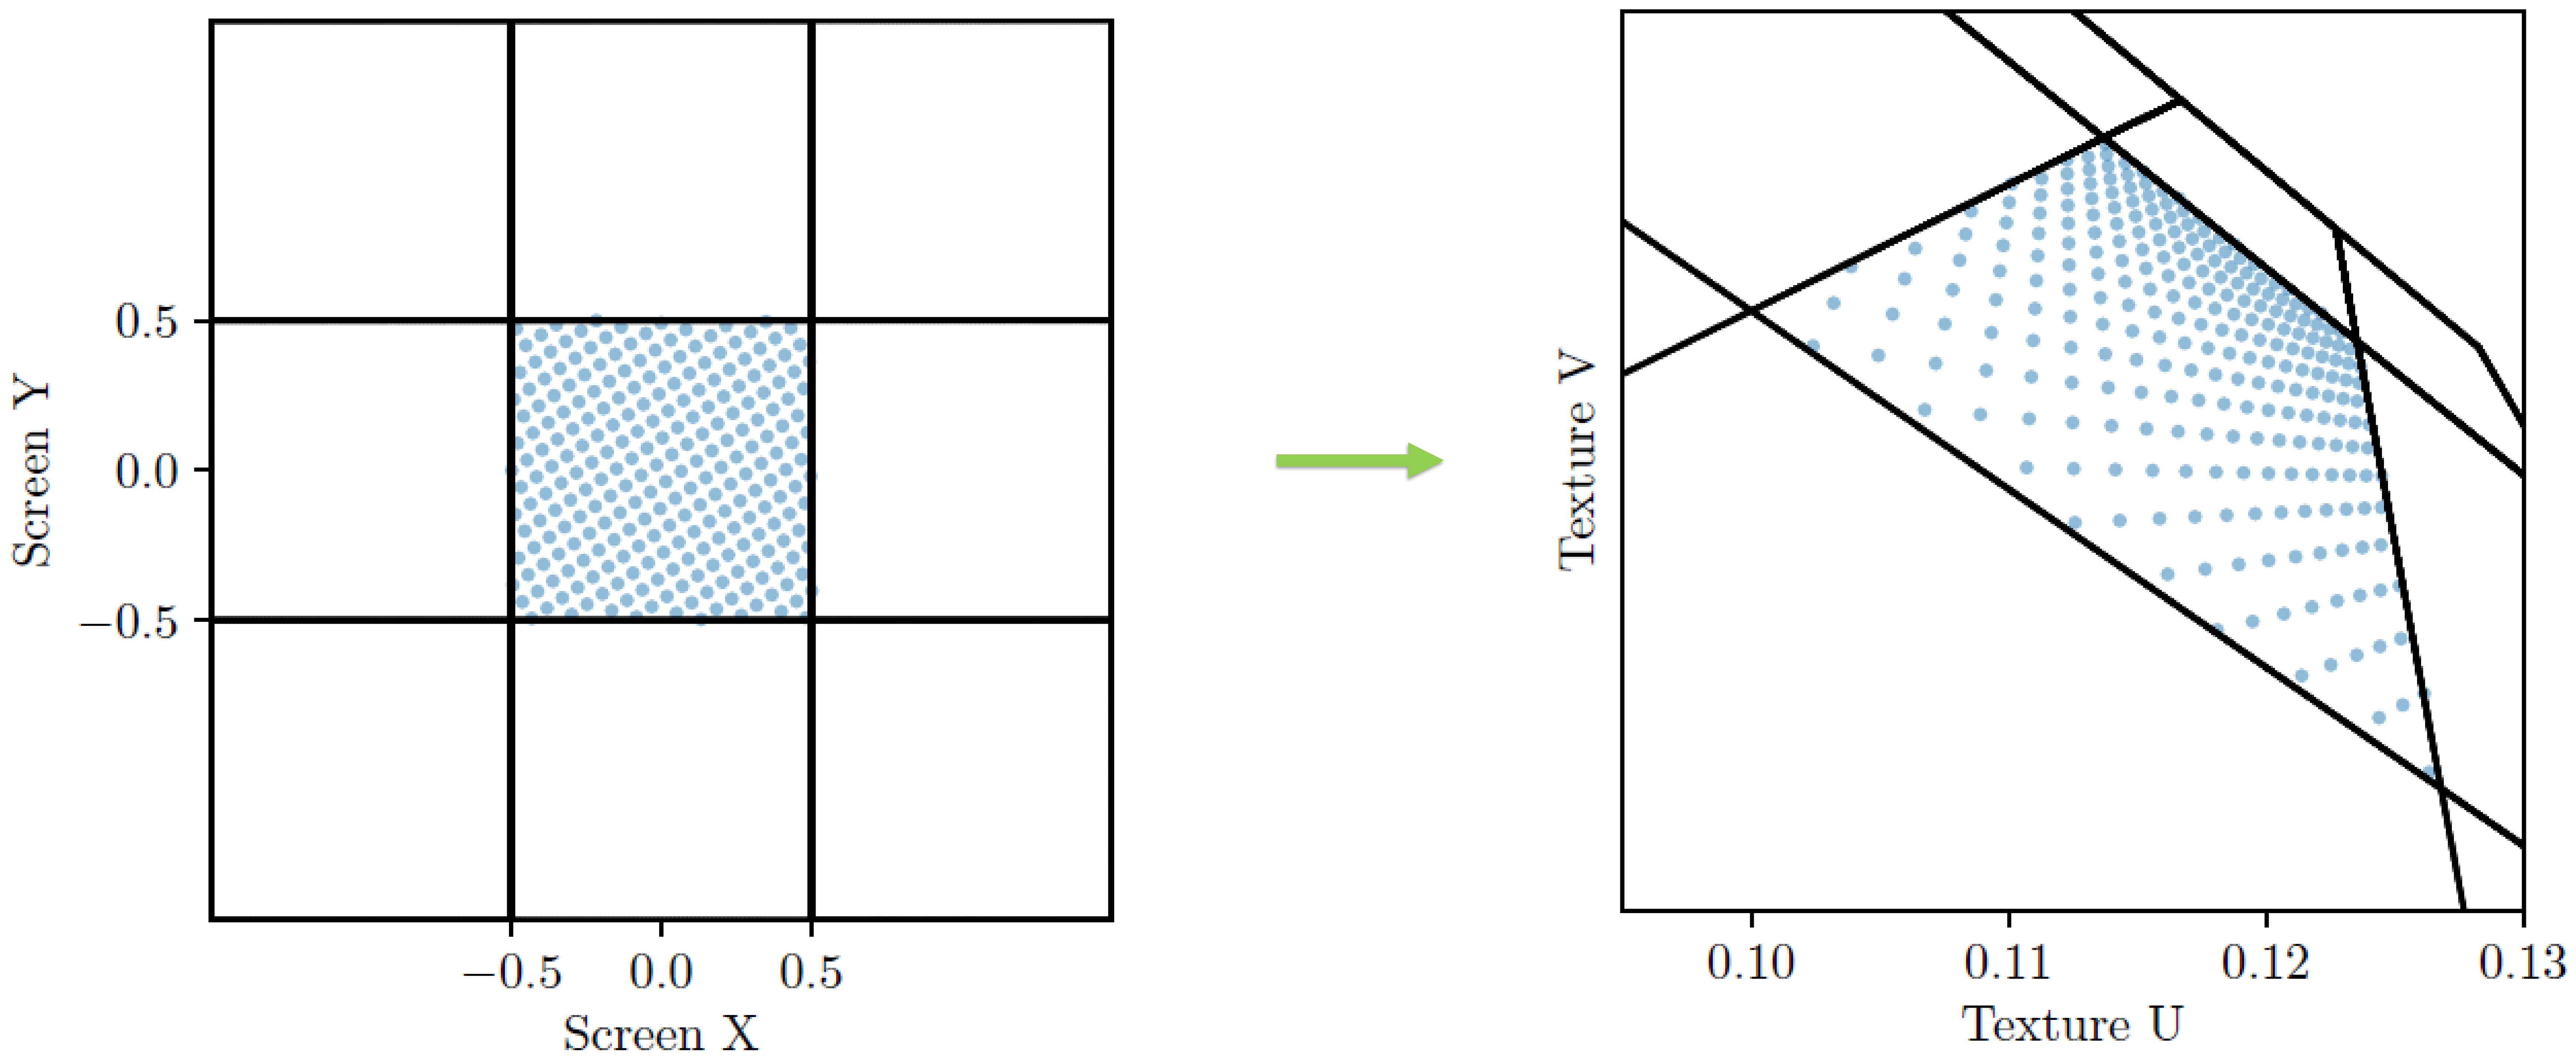

Minification
Stochastic Filtering After Shading

- Common practice – jitter the projection matrix for anti-aliasing reconstruction filter
- Used offline (e.g., MoonRay) and real-time (TAA, DLSS)
- Projects to trapezoid, minification supersampling -> filtering after shading!
- Add magnification/translation UV jitter -> **unified minification and magnification**

Minification vs magnification jitter

Appearance change and possible aliasing

Magnification specular appearance change

Filtering before shading

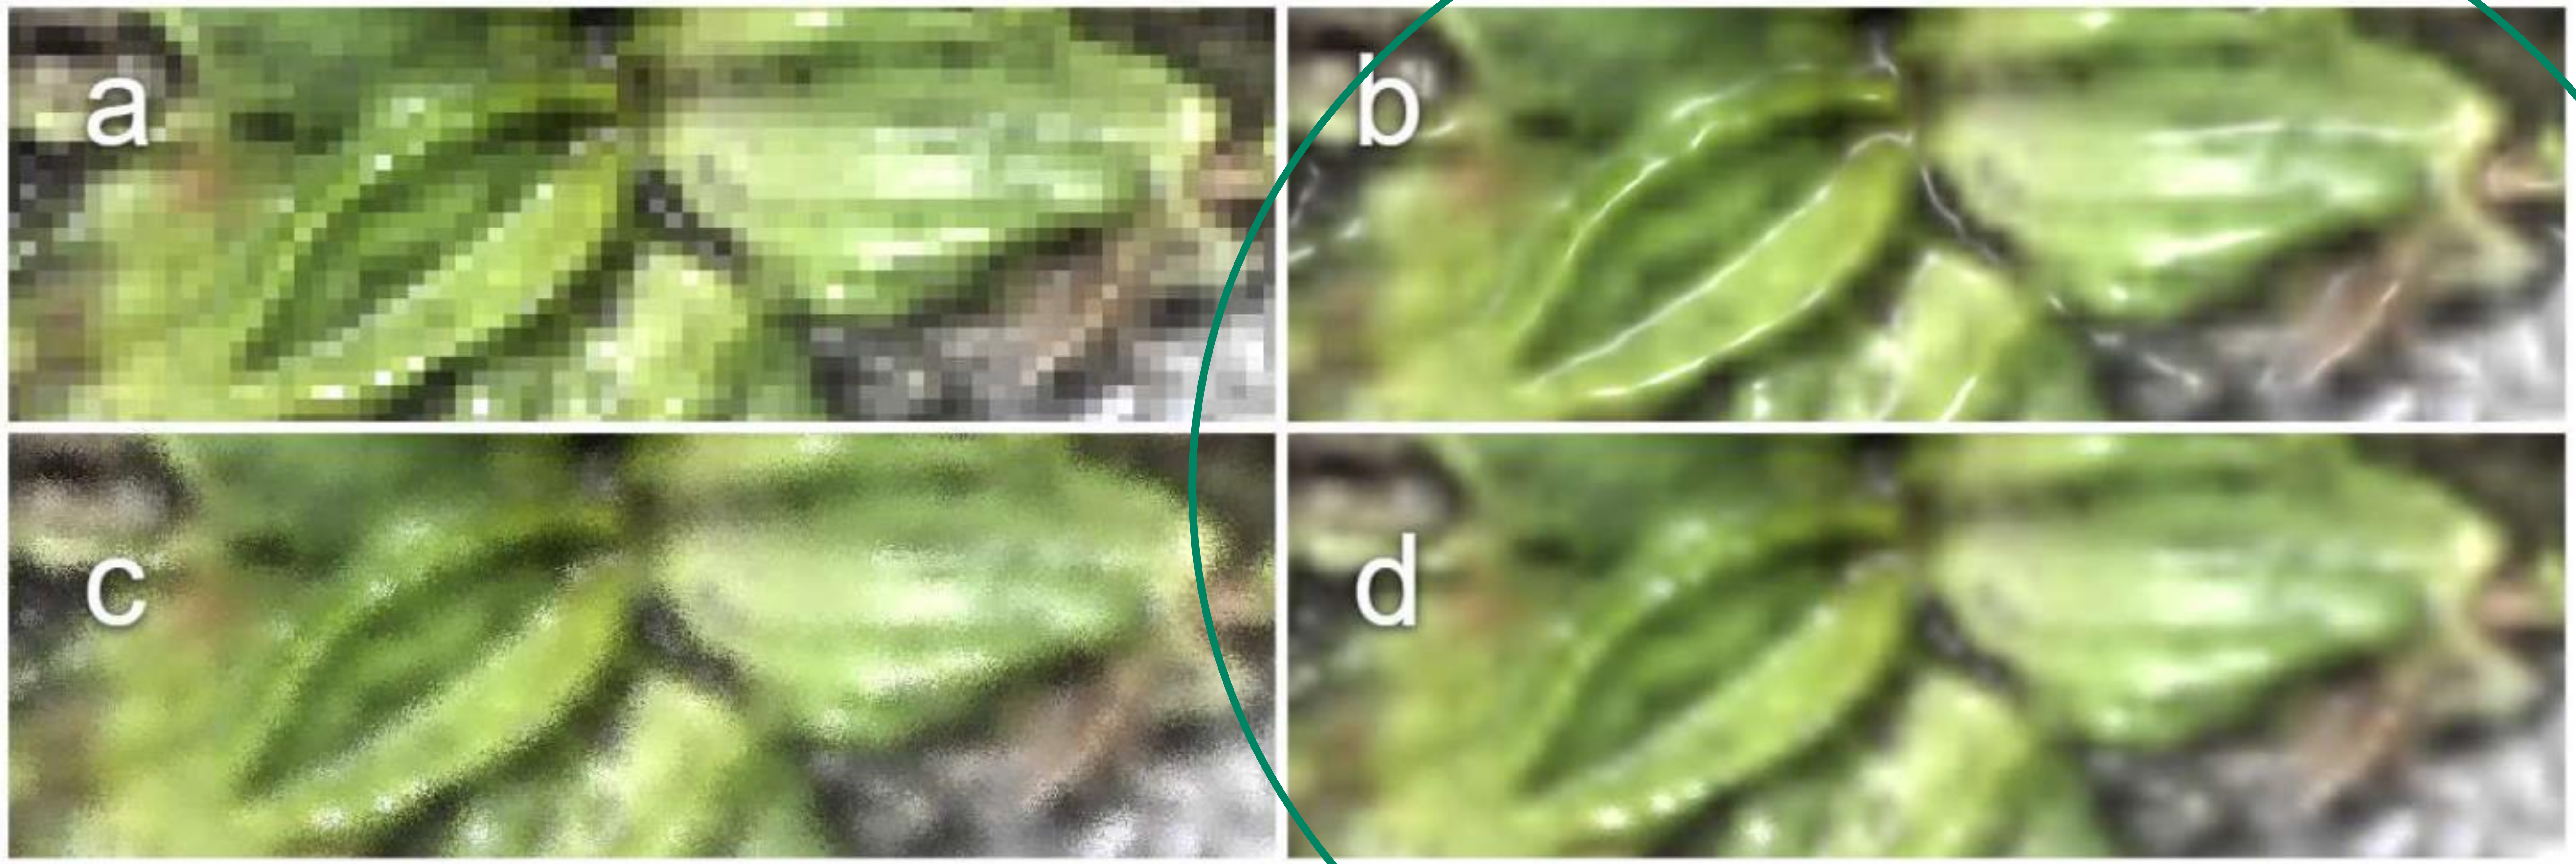

Filtering after shading

Appearance change explained

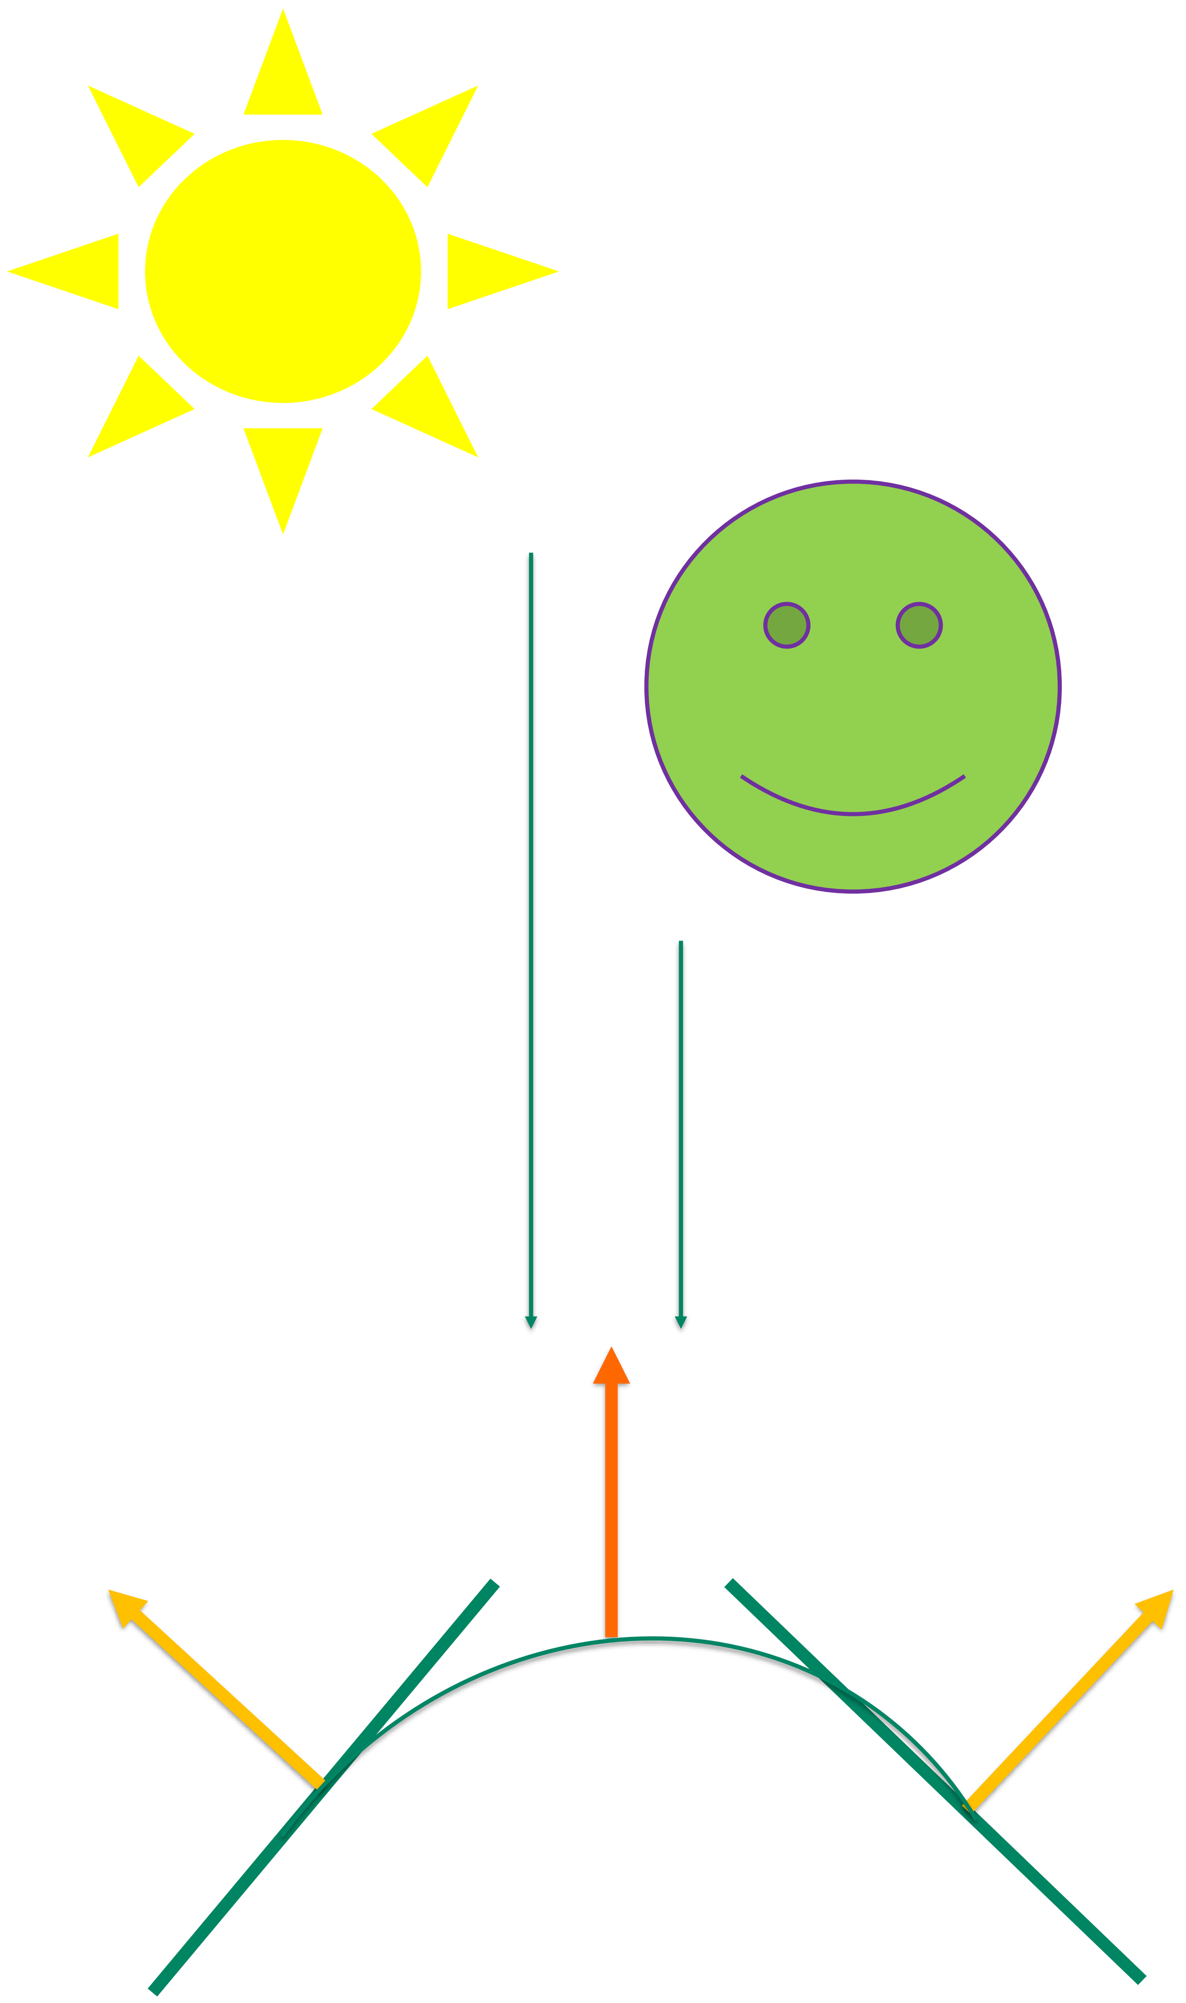

Filtering before shading:
Interpolated surface and normals

Appearance change explained

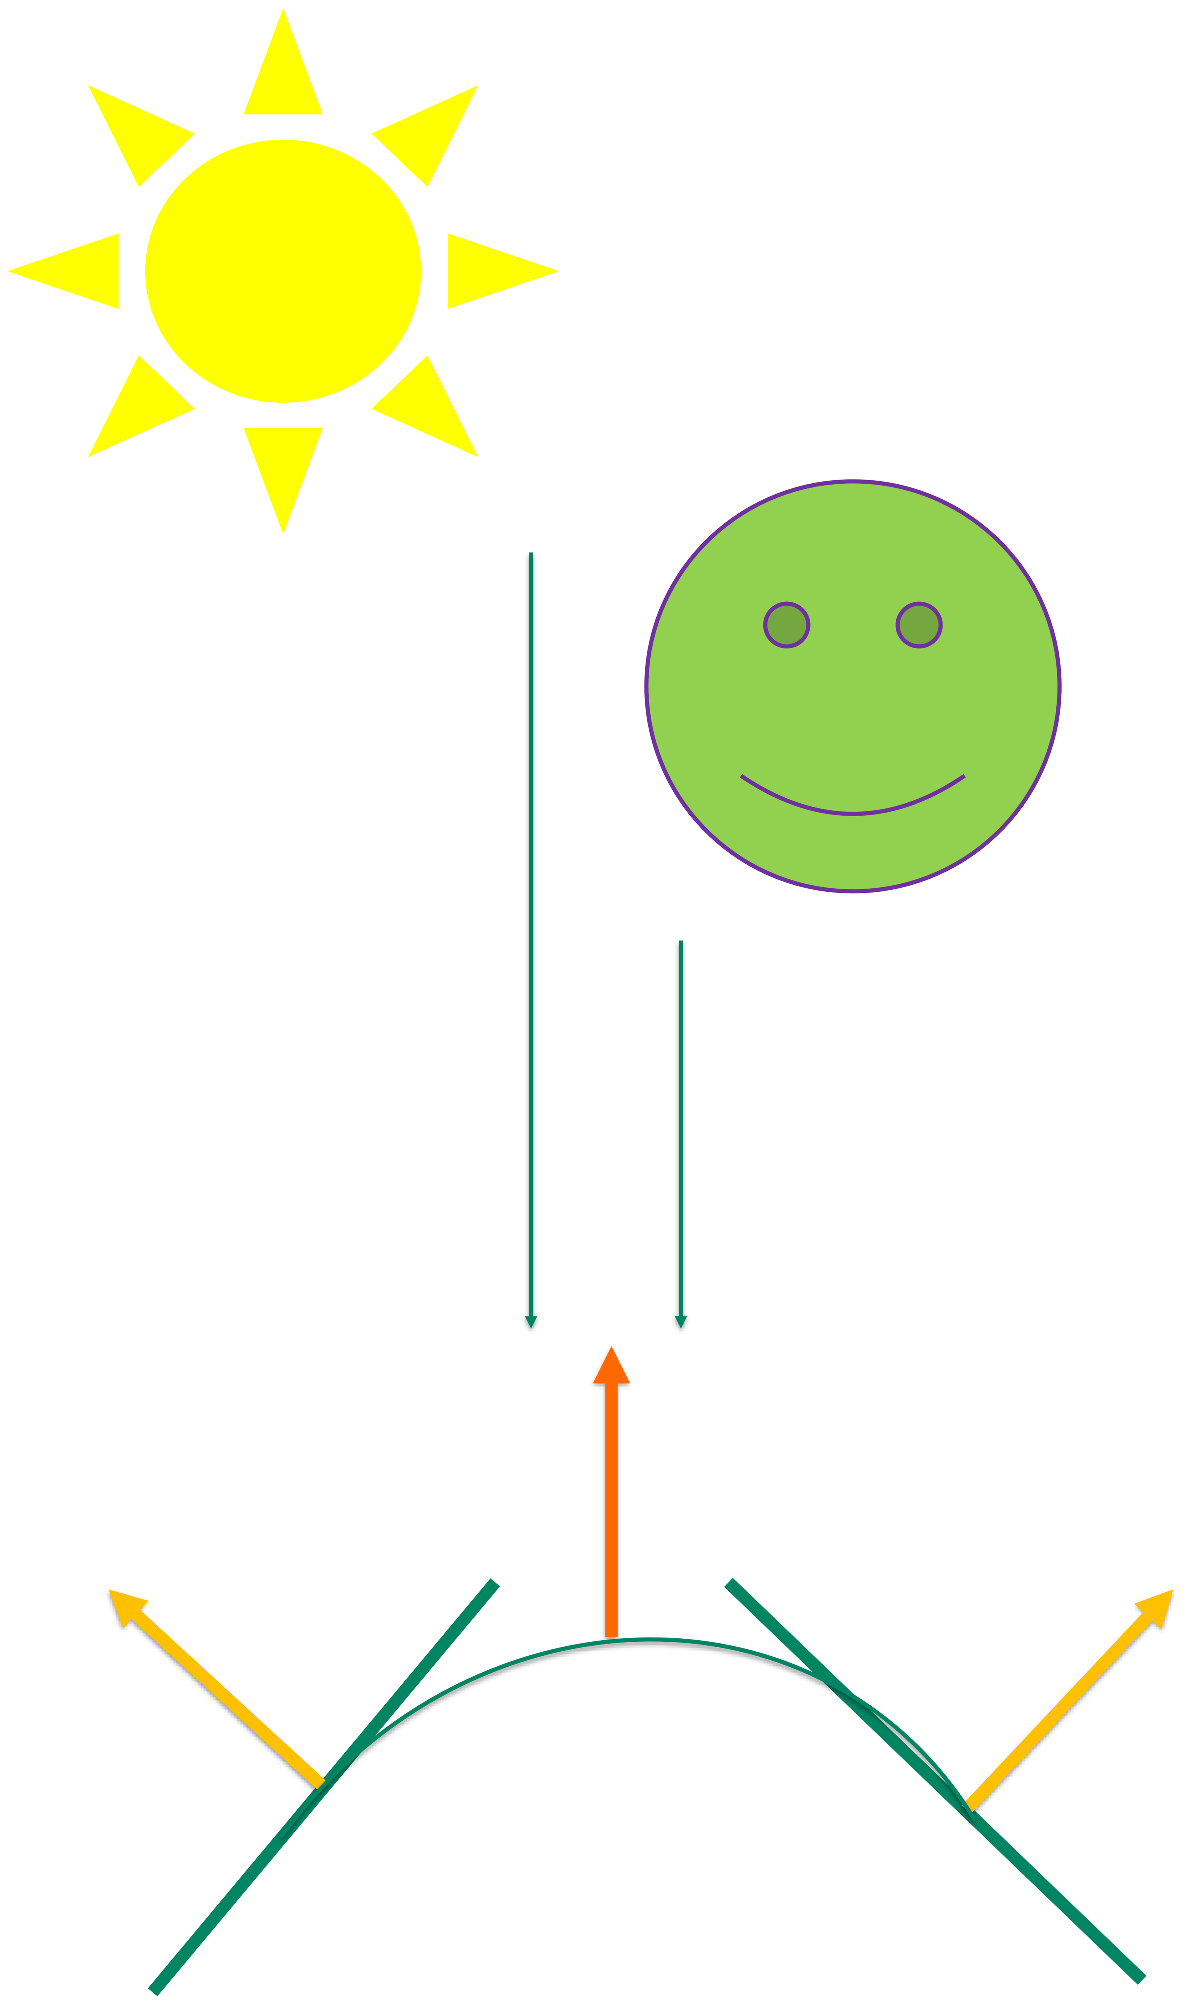

Filtering before shading:
Interpolated surface and normals

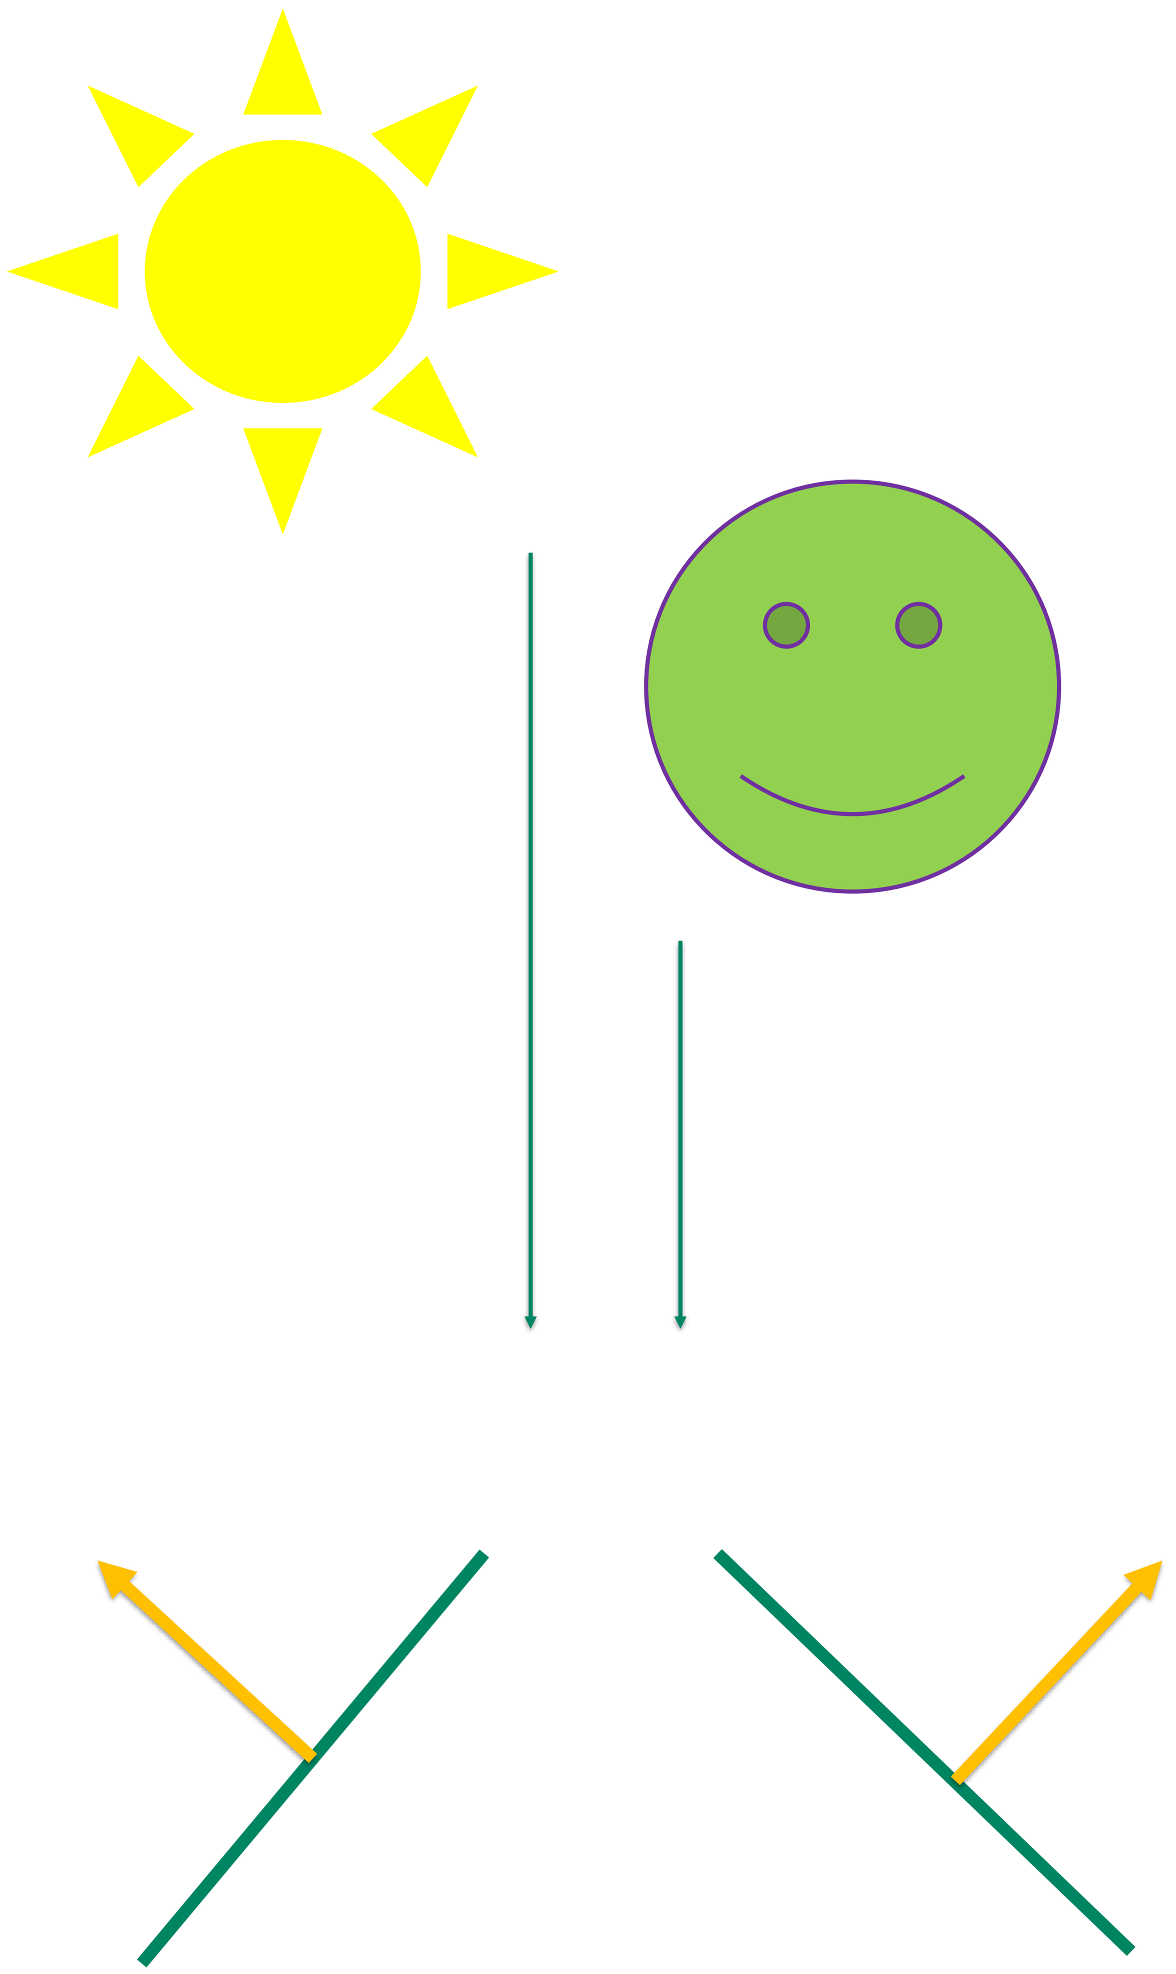

Filtering after shading:
Two adjacent geometric facets and normals

Filtering lighting does not produce surface curvature

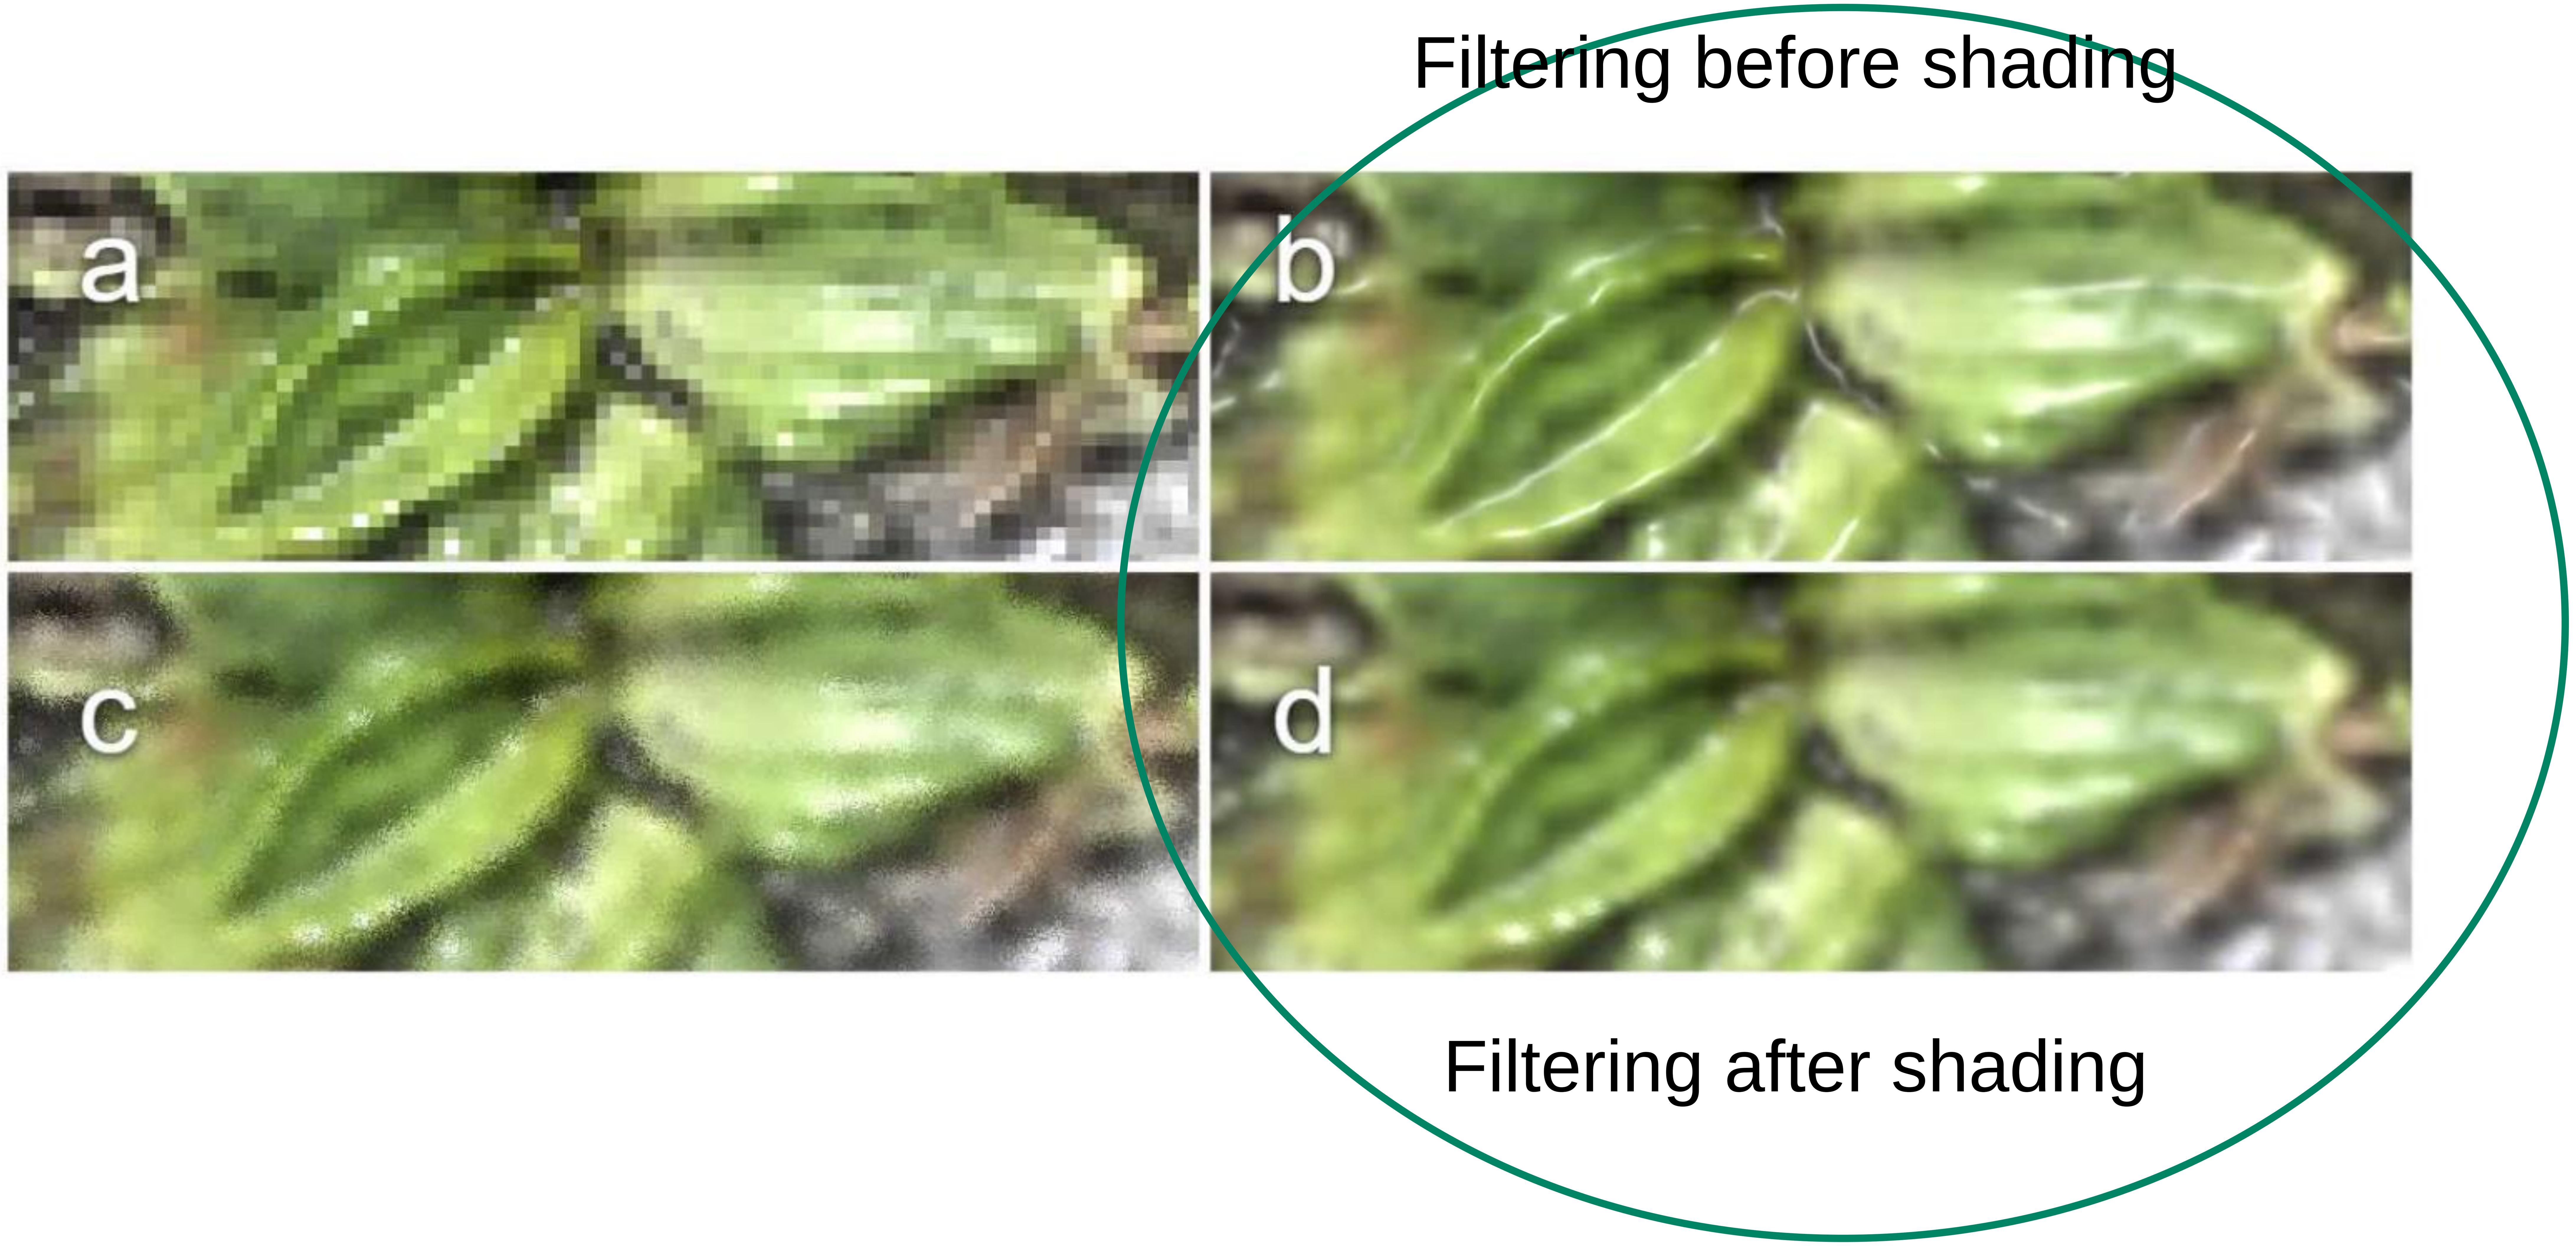

Filtering After Shading appearance difference

Note: it’s neither “good” or “bad”, depends on the intent

- But it changes the appearance – artists need to be aware!

Filtering After Shading appearance difference

Note: it's neither "good" or "bad", depends on the intent

- But it changes the appearance – artists need to be aware!

Our original assumption:

- **"Textures are authored by artists with ~1-1 pixel-texel ratio"**
- Sometimes can be violated! Relying on smooth curvature

Filtering After Shading appearance difference

Note: it's neither "good" or "bad", depends on the intent

- But it changes the appearance – artists need to be aware!

Our original assumption:

- **"Textures are authored by artists with ~1-1 pixel-texel ratio"**
- Sometimes can be violated! Relying on smooth curvature
- Extreme example of relying on interpolation: SDF fonts

(a) 64x64 texture, alpha-blended (b) 64x64 texture, alpha tested (c) 64x64 texture using our technique

***Improved Alpha-Tested Magnification for Vector Textures and Special Effects*, Chris Green, Siggraph 2007**

Worse example – magnification aliasing

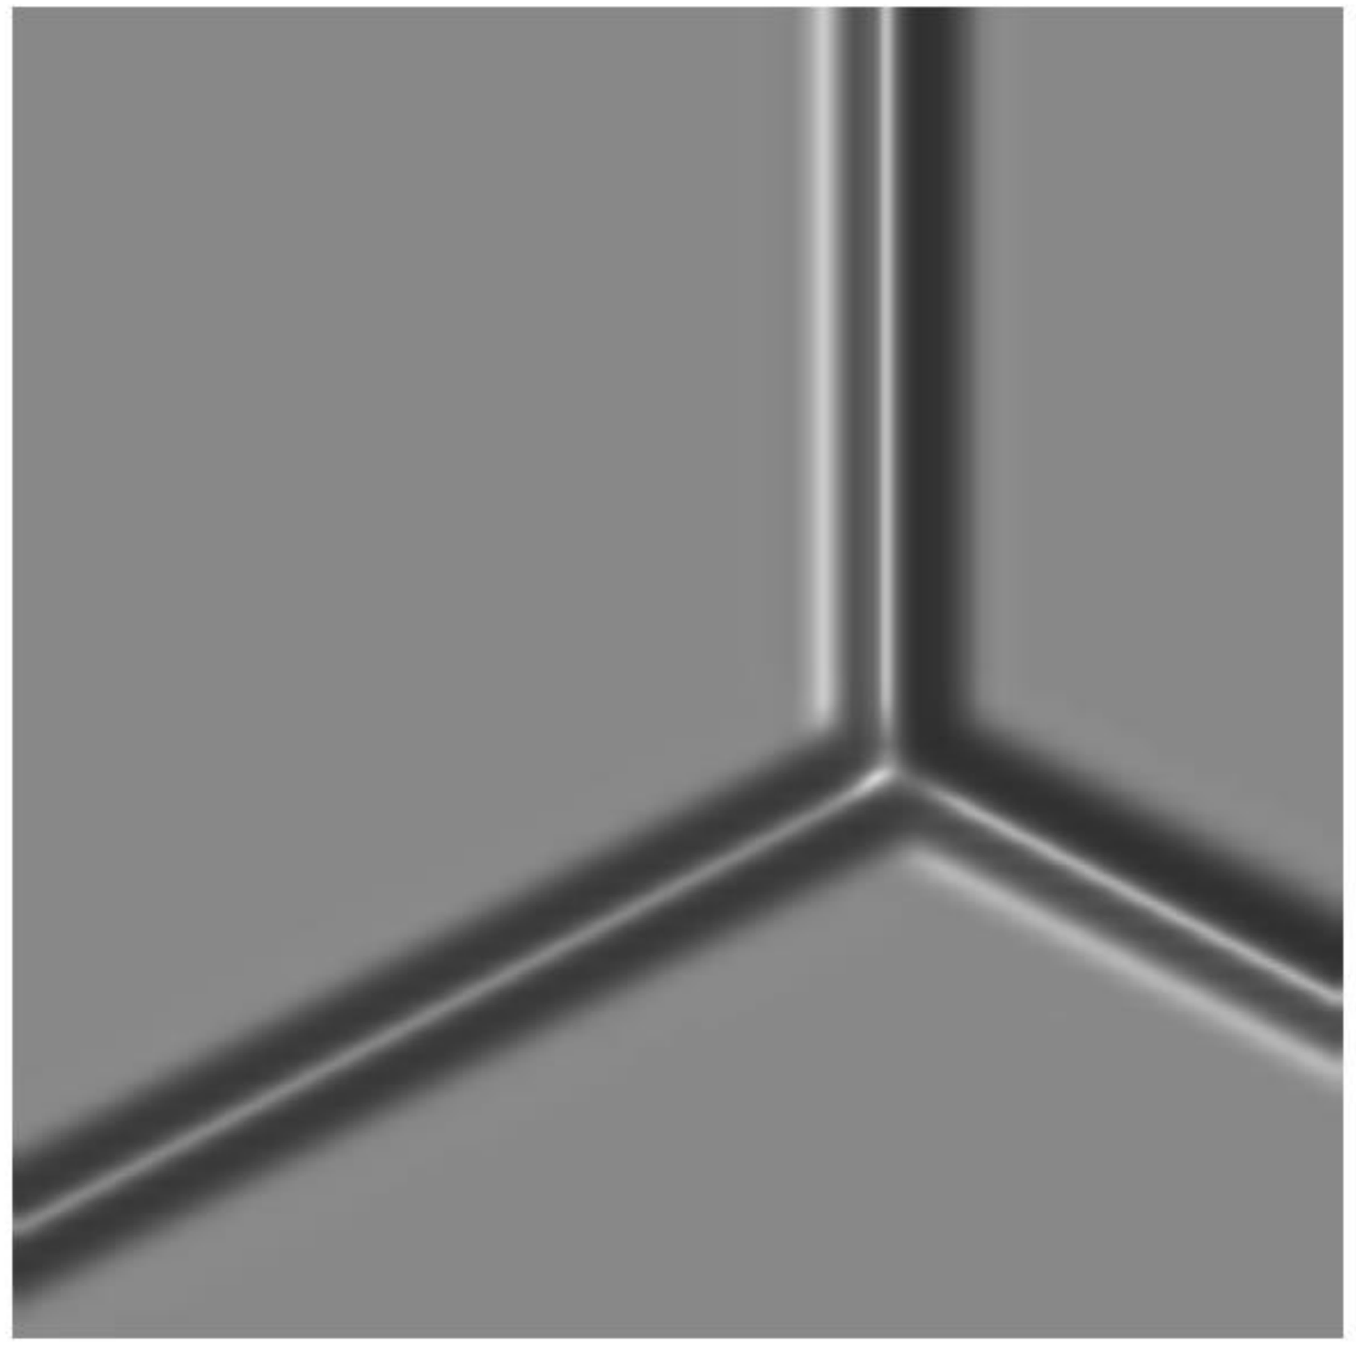

Filtering before shading

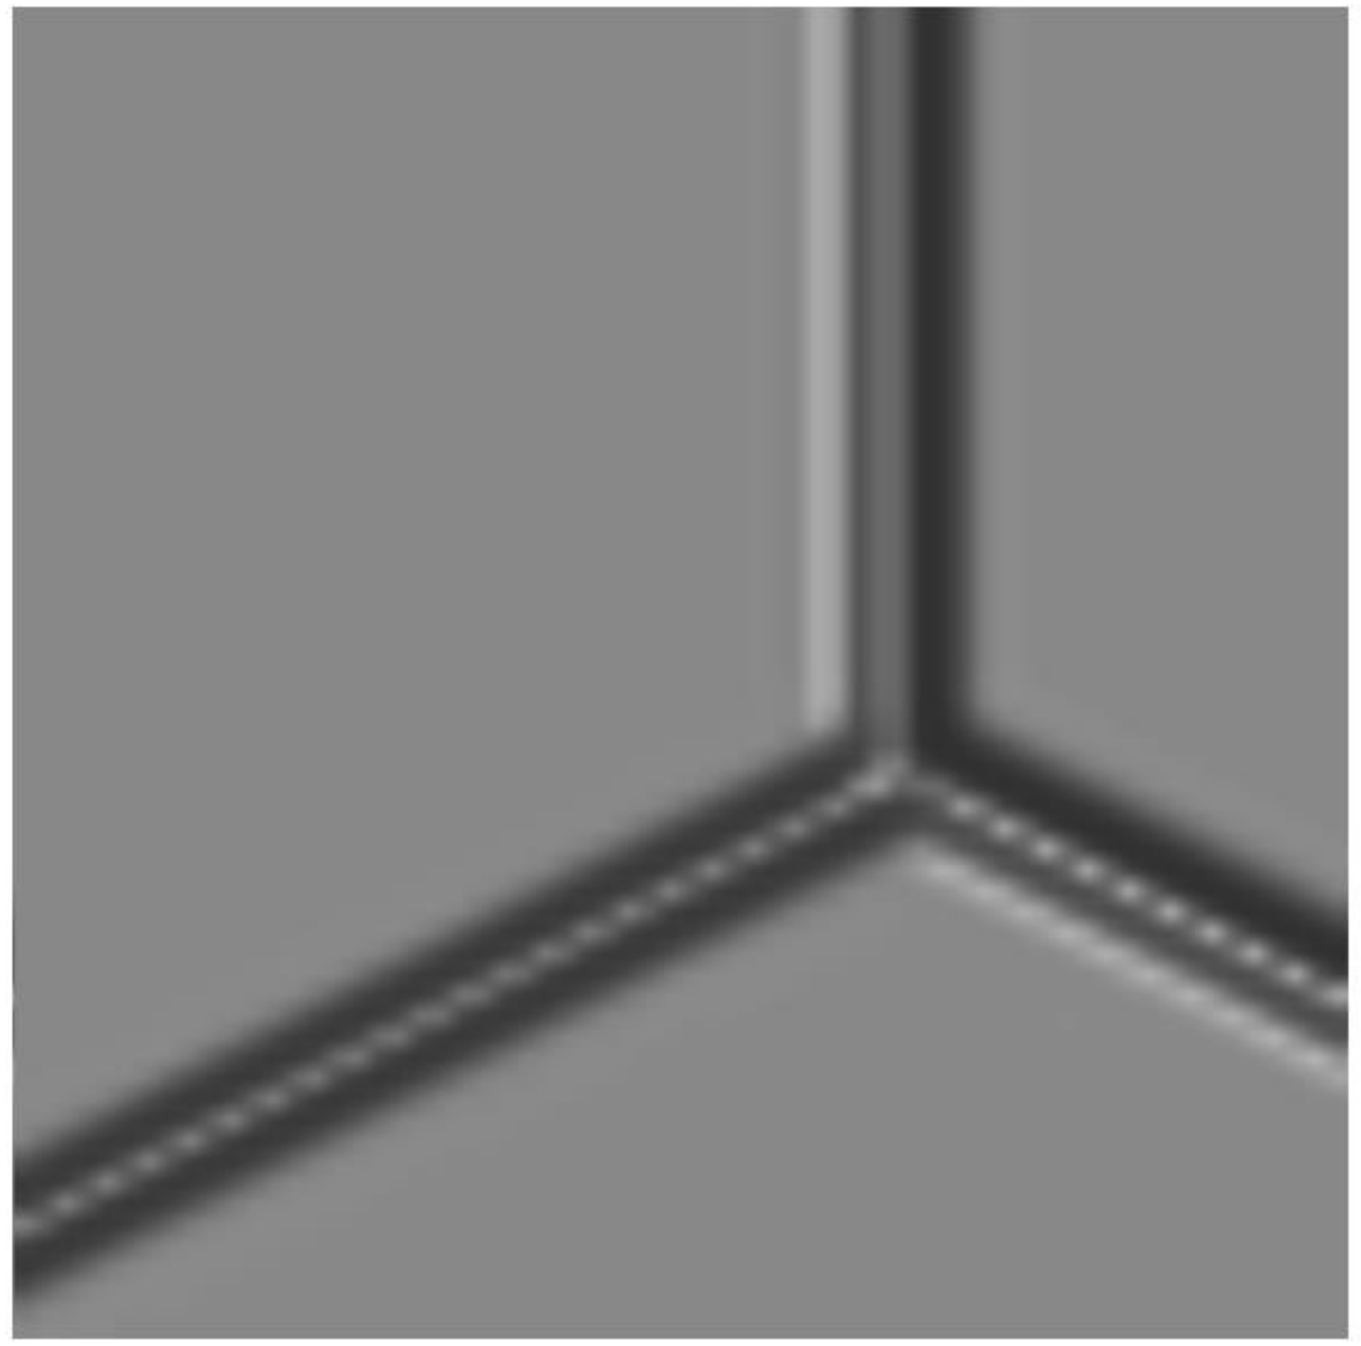

Filtering after shading

Non-linearity introduced aliasing

For formal analysis, please see the paper

- Any non-linearity always introduces new, higher signal frequencies (“harmonics”)
- sin(x)^2 == (sin(2x)+1)/2

Non-linearity introduced aliasing

For formal analysis, please see the paper

- Any non-linearity always introduces new, higher signal frequencies (“harmonics”)
- sin(x)^2 == (sin(2x)+1)/2
- The same for further powers, sums of sums generate sum of frequencies (“*intermodulation*”)
- This is an extreme example, but **every** non-trivial non-linearity introduces higher frequencies
- (Can analyze through Taylor expansion etc.)

Non-linearity introduced aliasing

When applied to discrete signals…

- **Those frequencies alias immediately**
- Amount of aliasing depends on the non-linearity, original spectral content, phases
- **We did not see this problem for many months!**

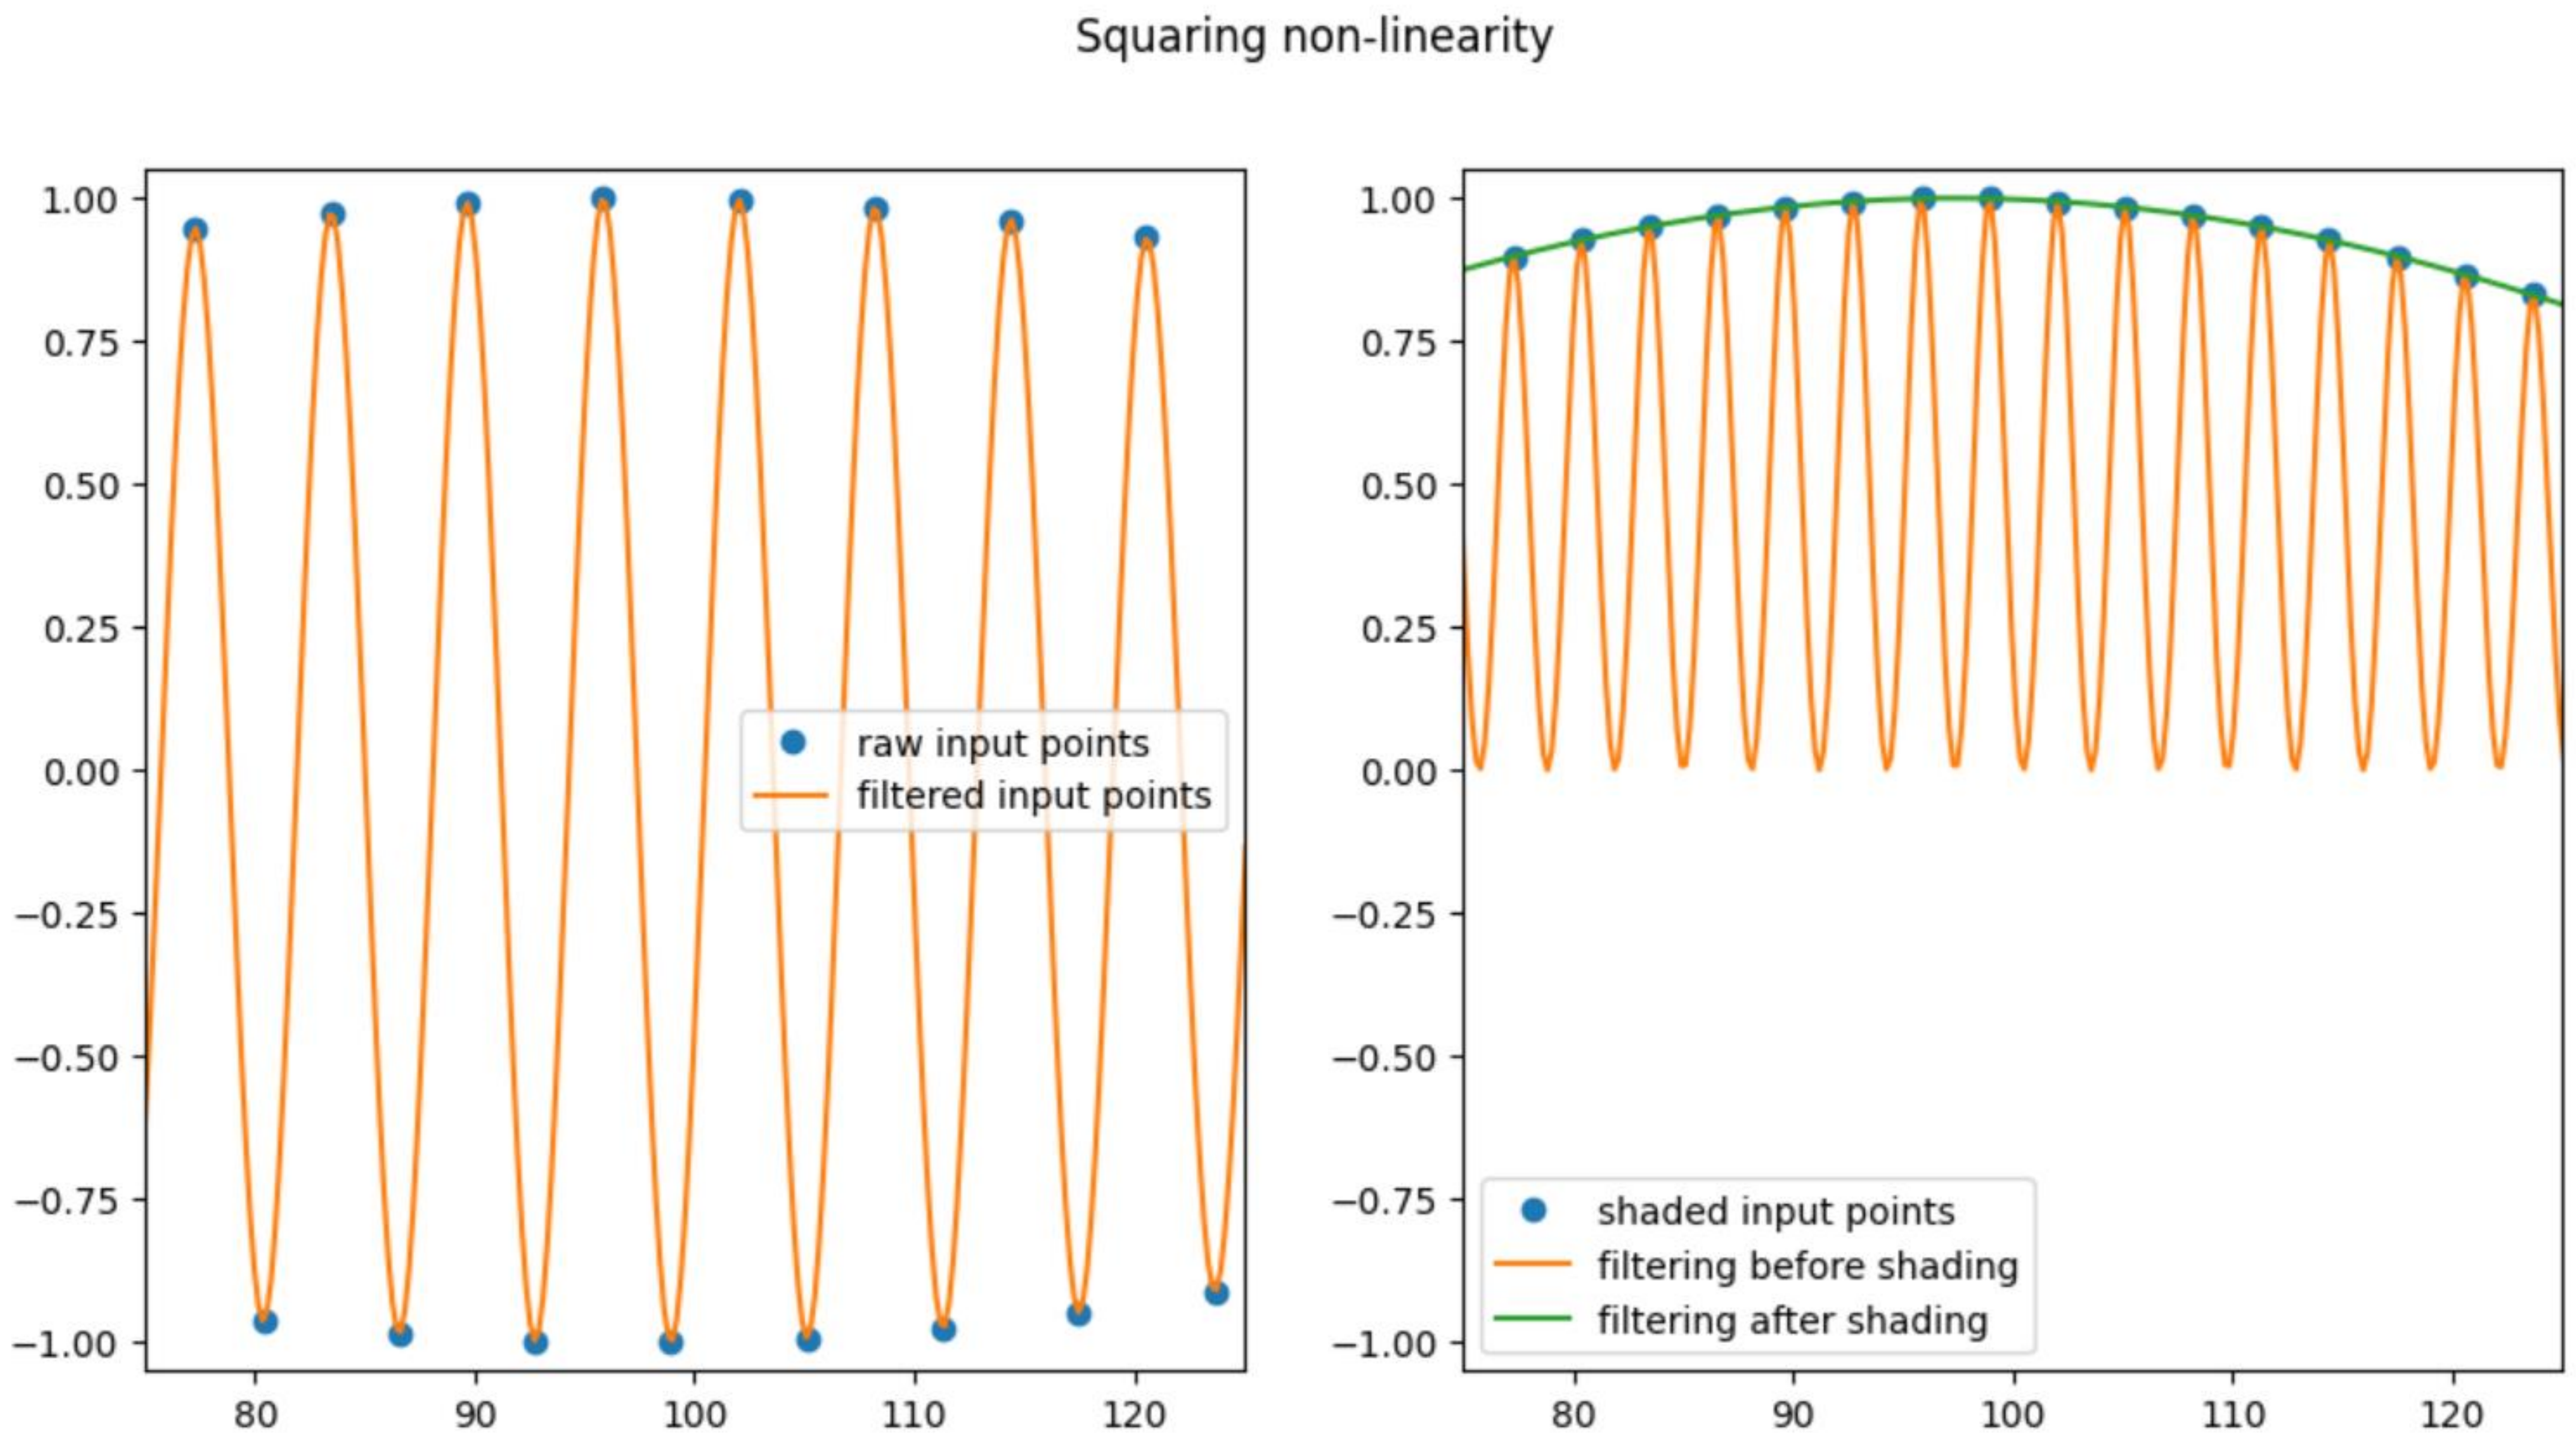

Specular-like scenario

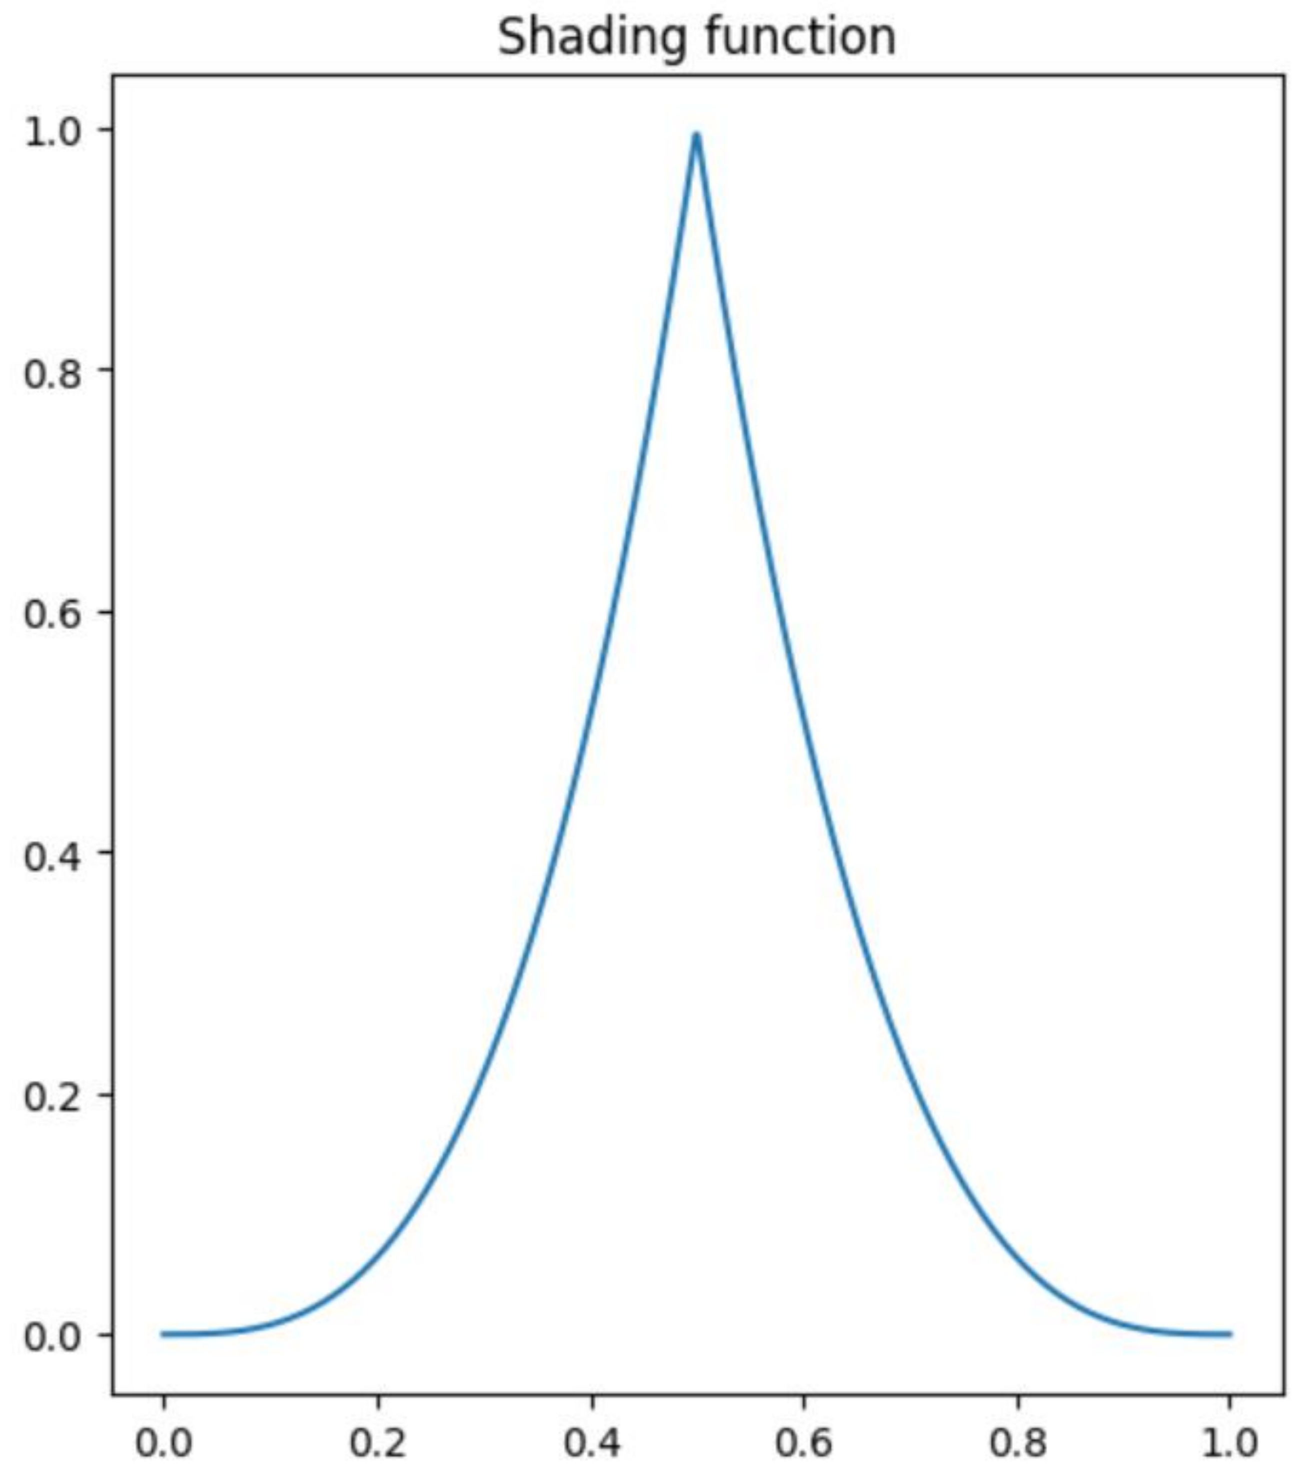

Specular-like scenario

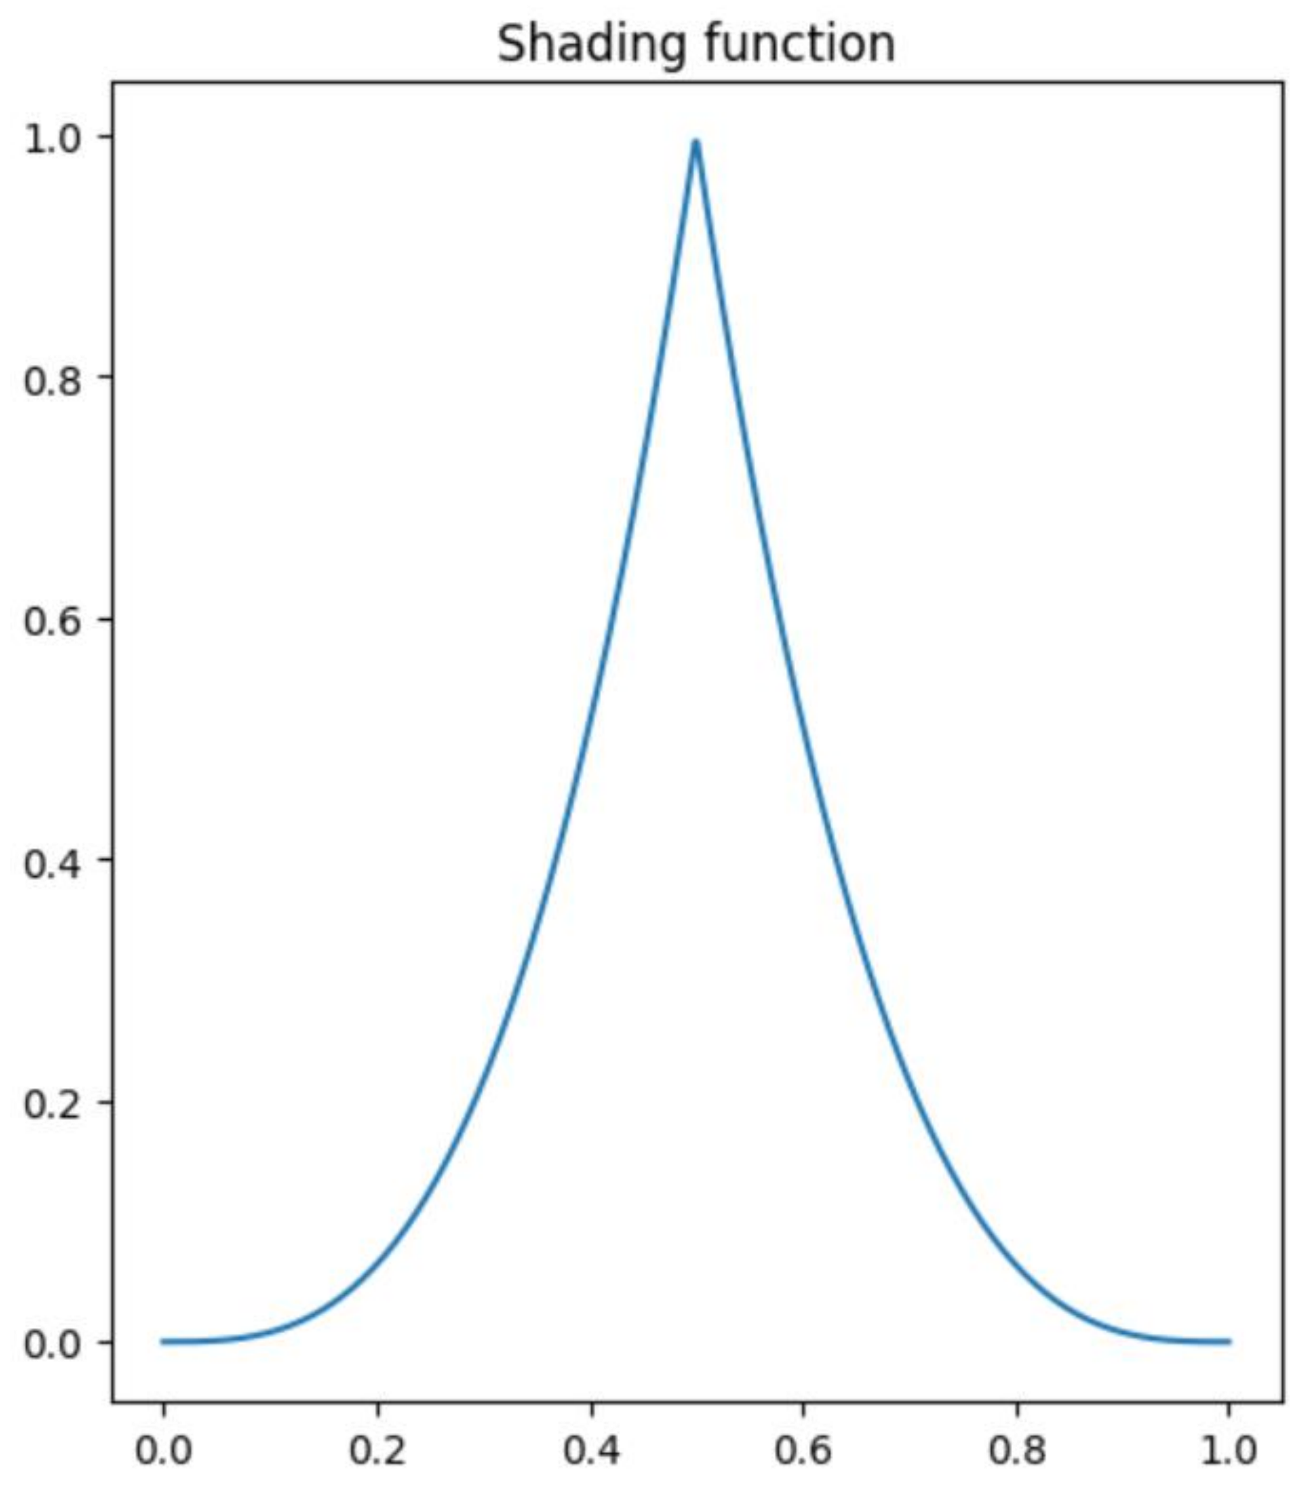

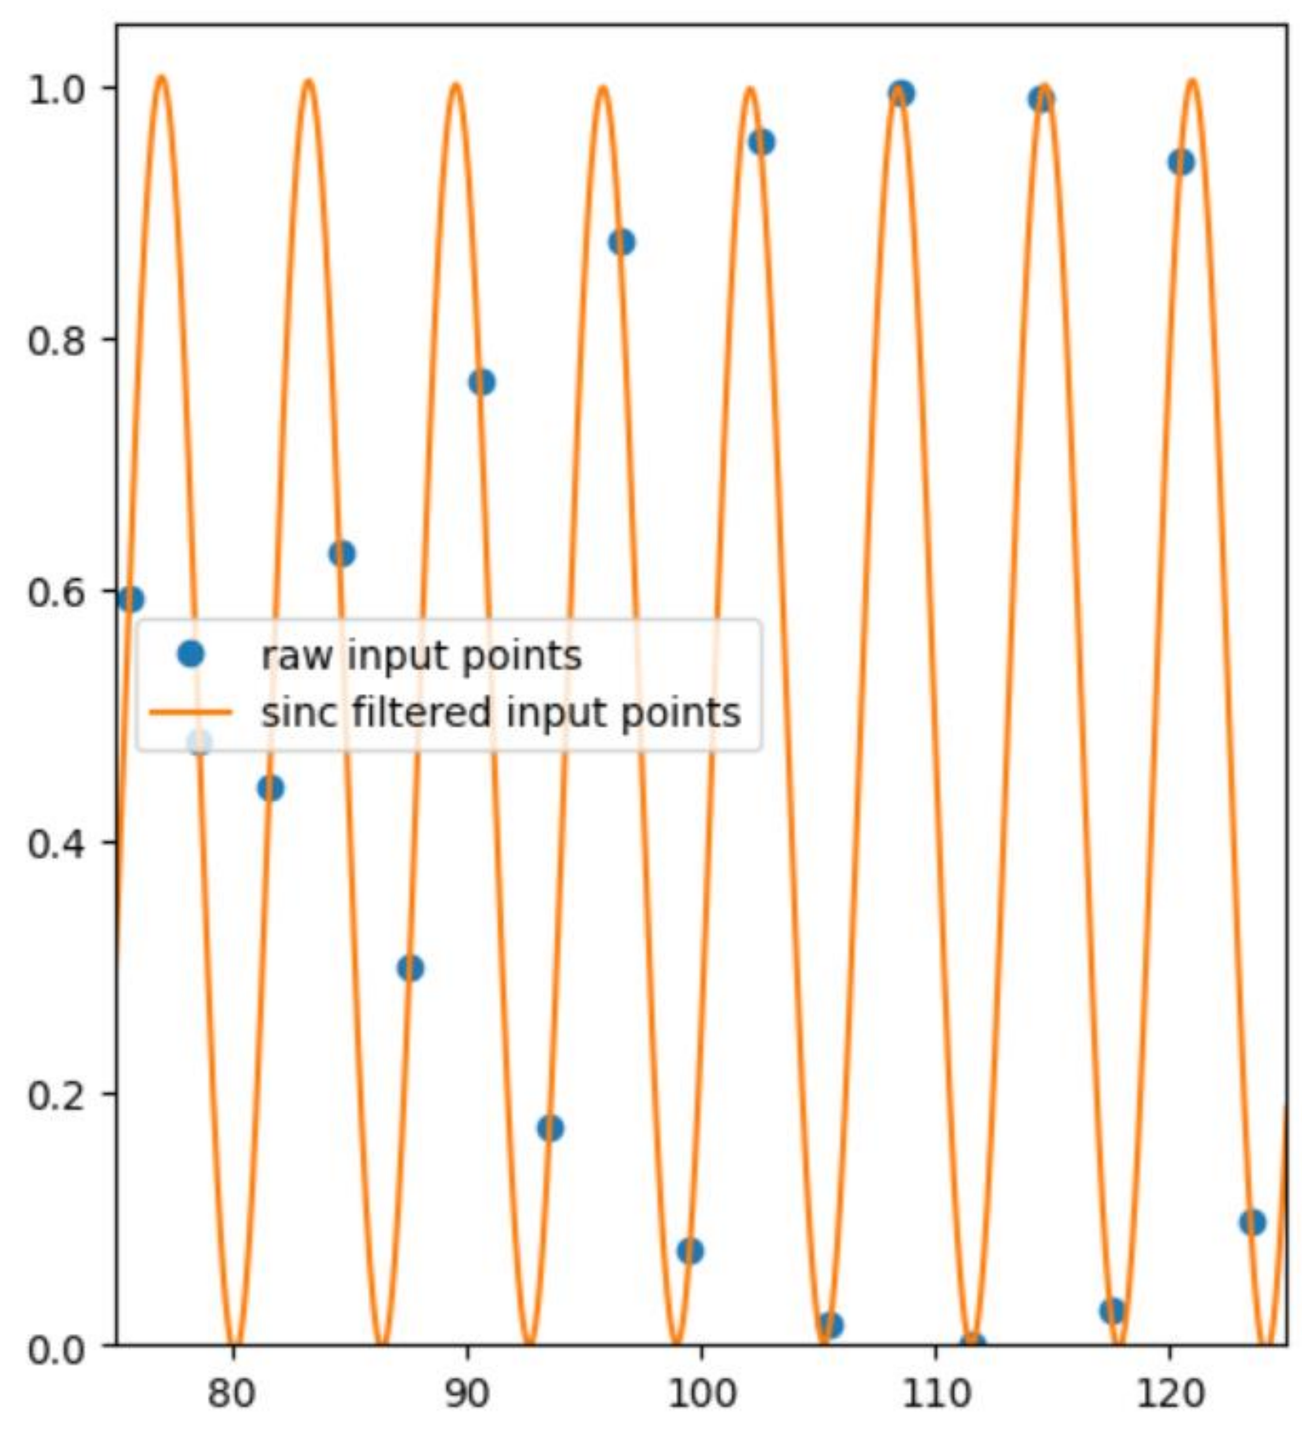

Specular-like scenario

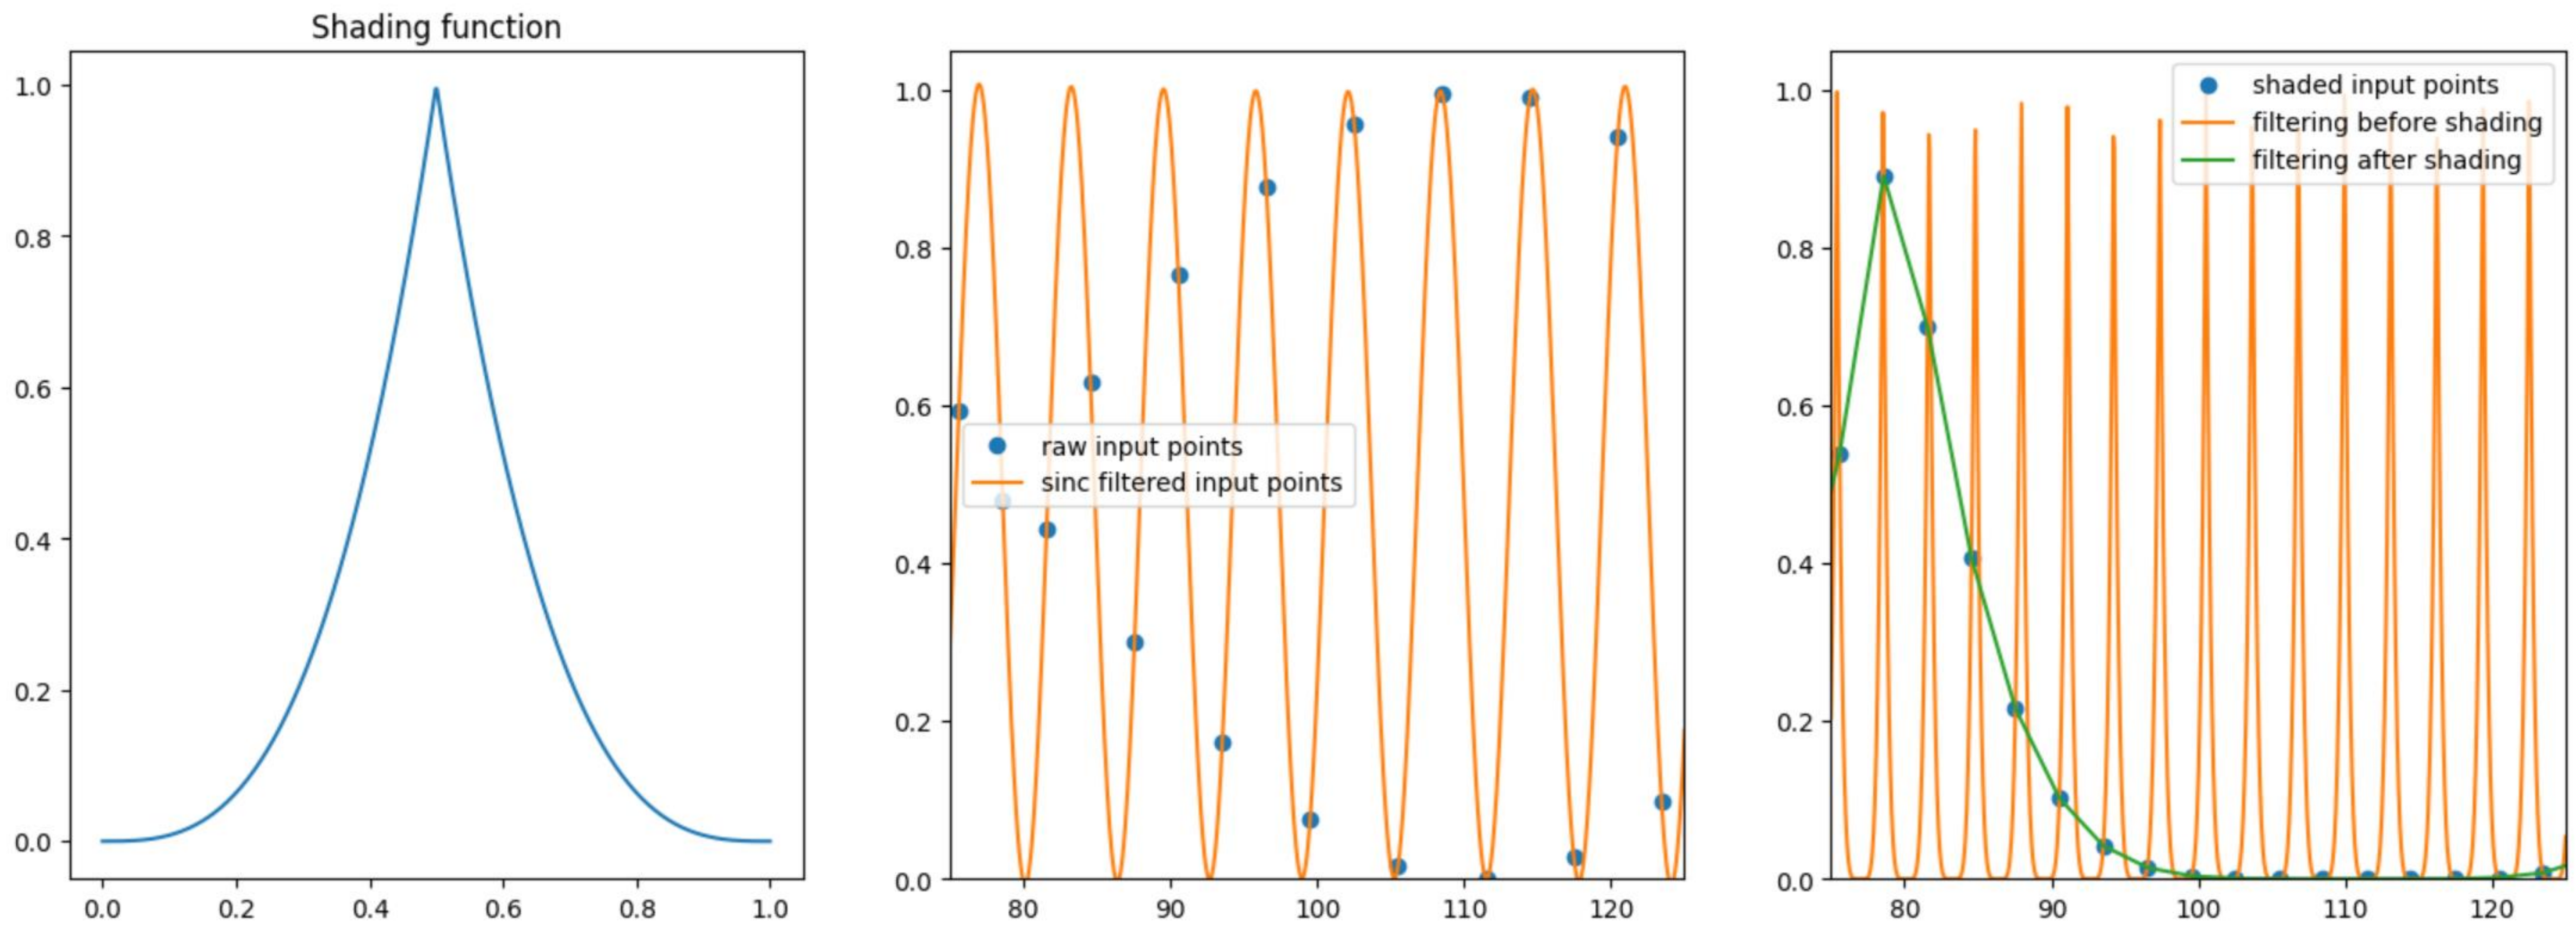

Non-linearity introduced aliasing

Magnification: screen Nyquist higher than texture Nyquist

- After magnification, more bandwidth headroom before aliasing happens
- If we apply non-linearity first, we alias more and cannot recover

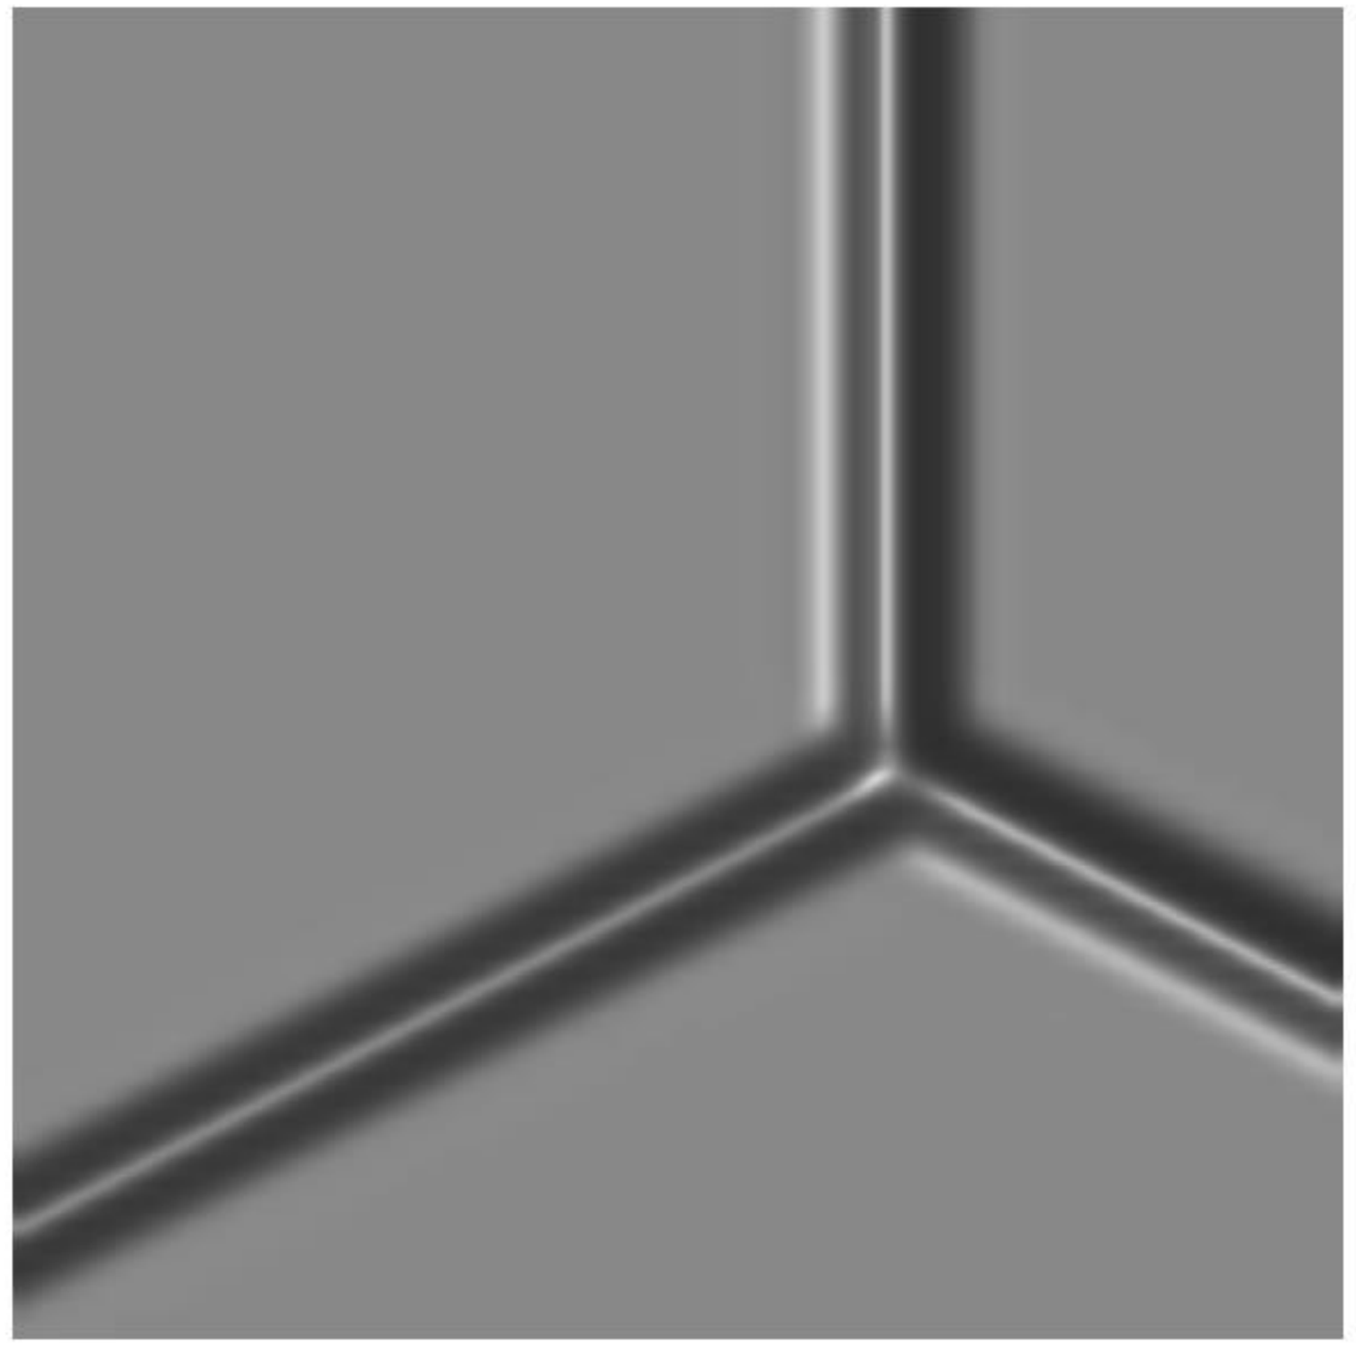

Filtering before shading

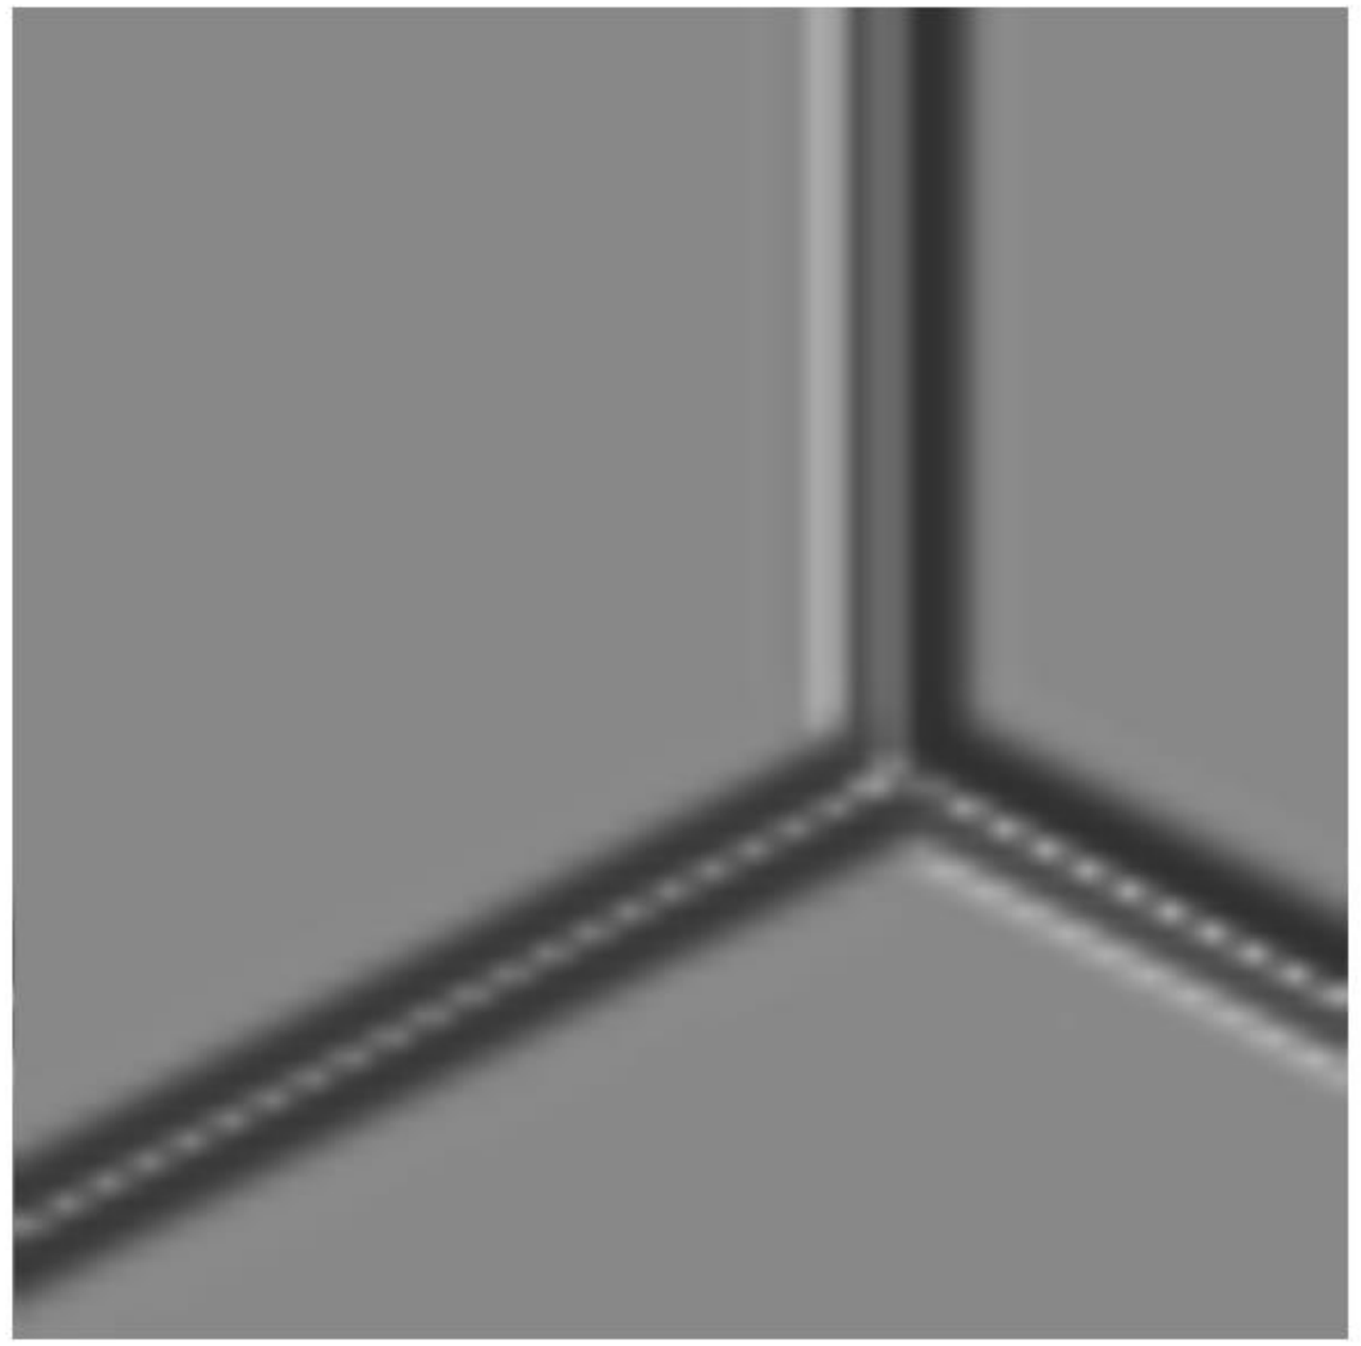

Filtering after shading

Non-linearity introduced aliasing

Minification: screen Nyquist lower than texture Nyquist

- More headroom -> less aliasing
- Filtering after shading can remove most of the nonlinearity-induced aliasing!

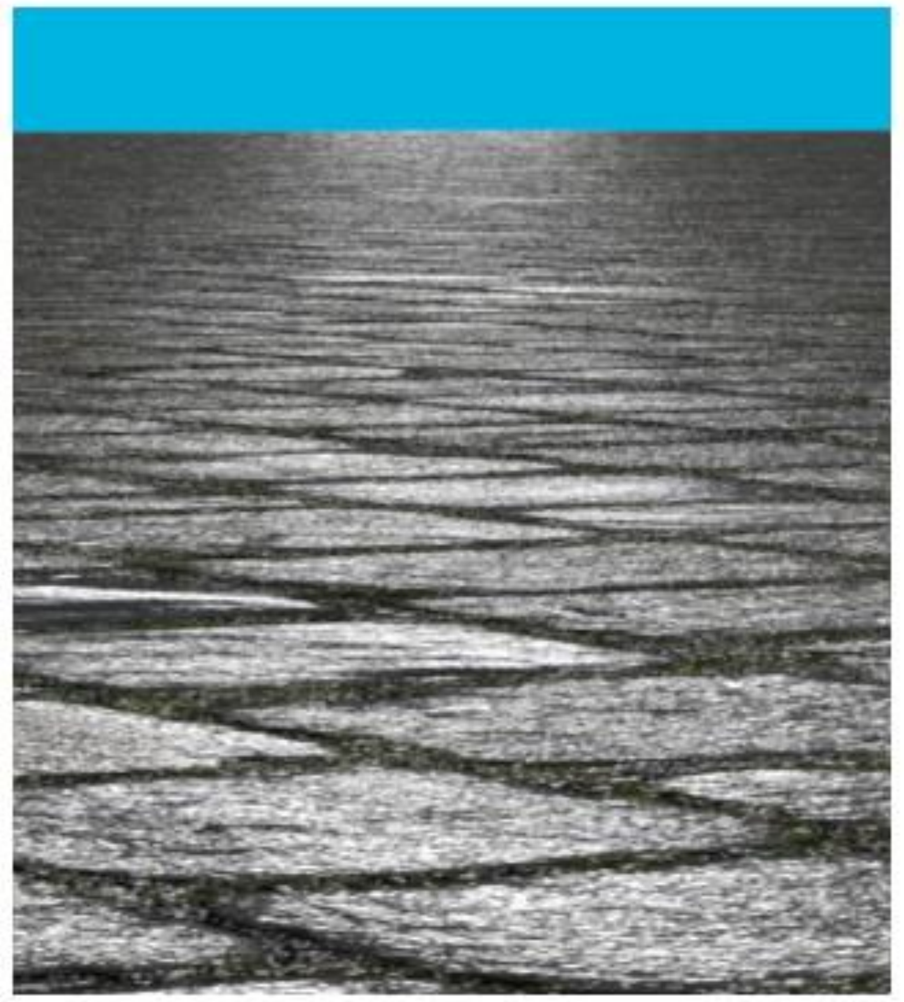

Filter after shading (real-time implementation)

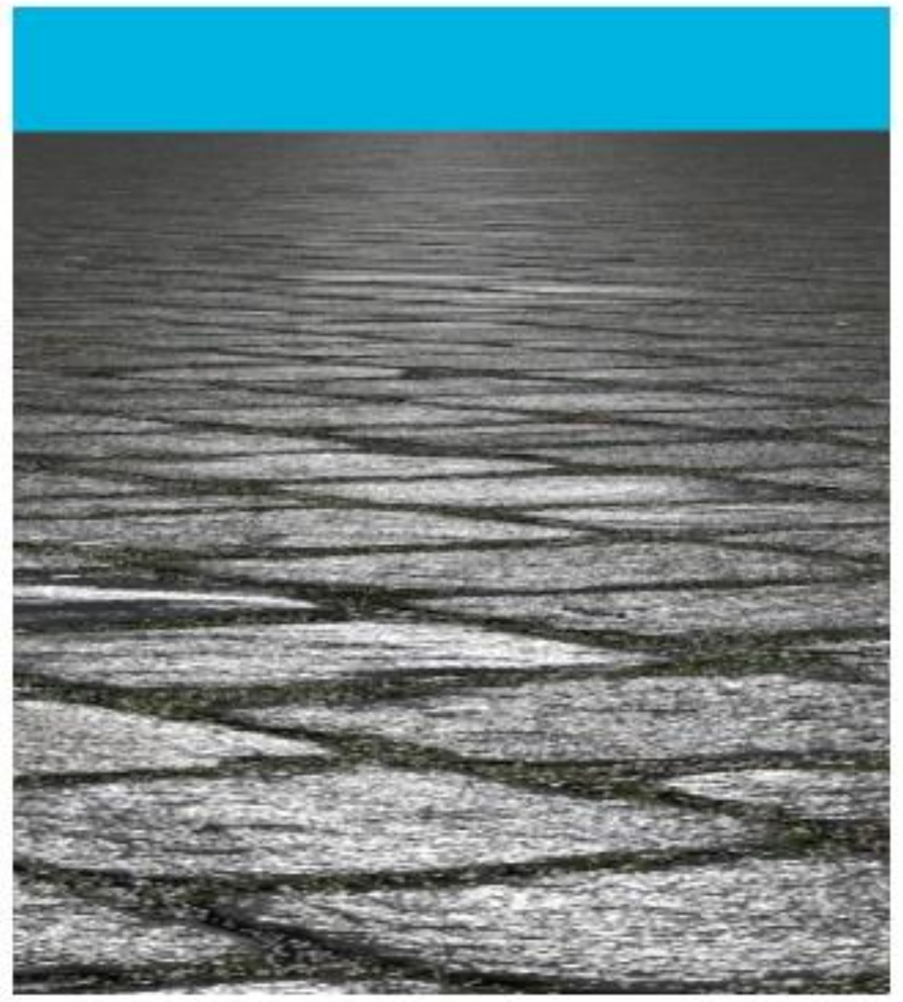

Reference

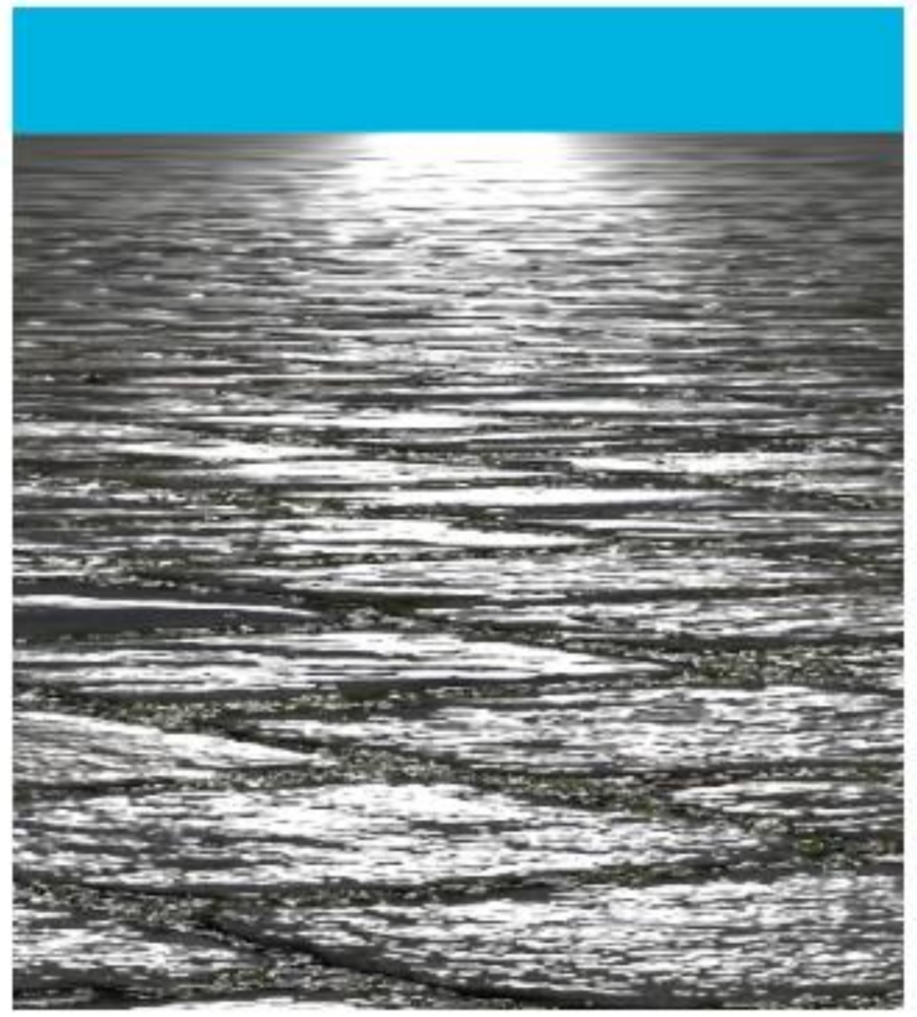

Filter before shading

Results

Appearance preservation – real time

(a) Filter Before 1 spp

(b) Filter Before 1 spp + DLSS

(c) Filter Before 1024 spp

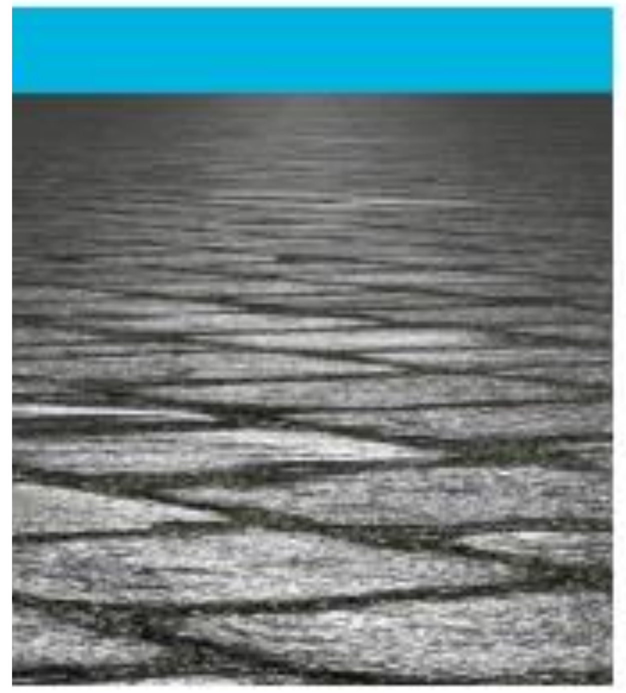

(h) Reference 1024 spp

Appearance preservation – real time

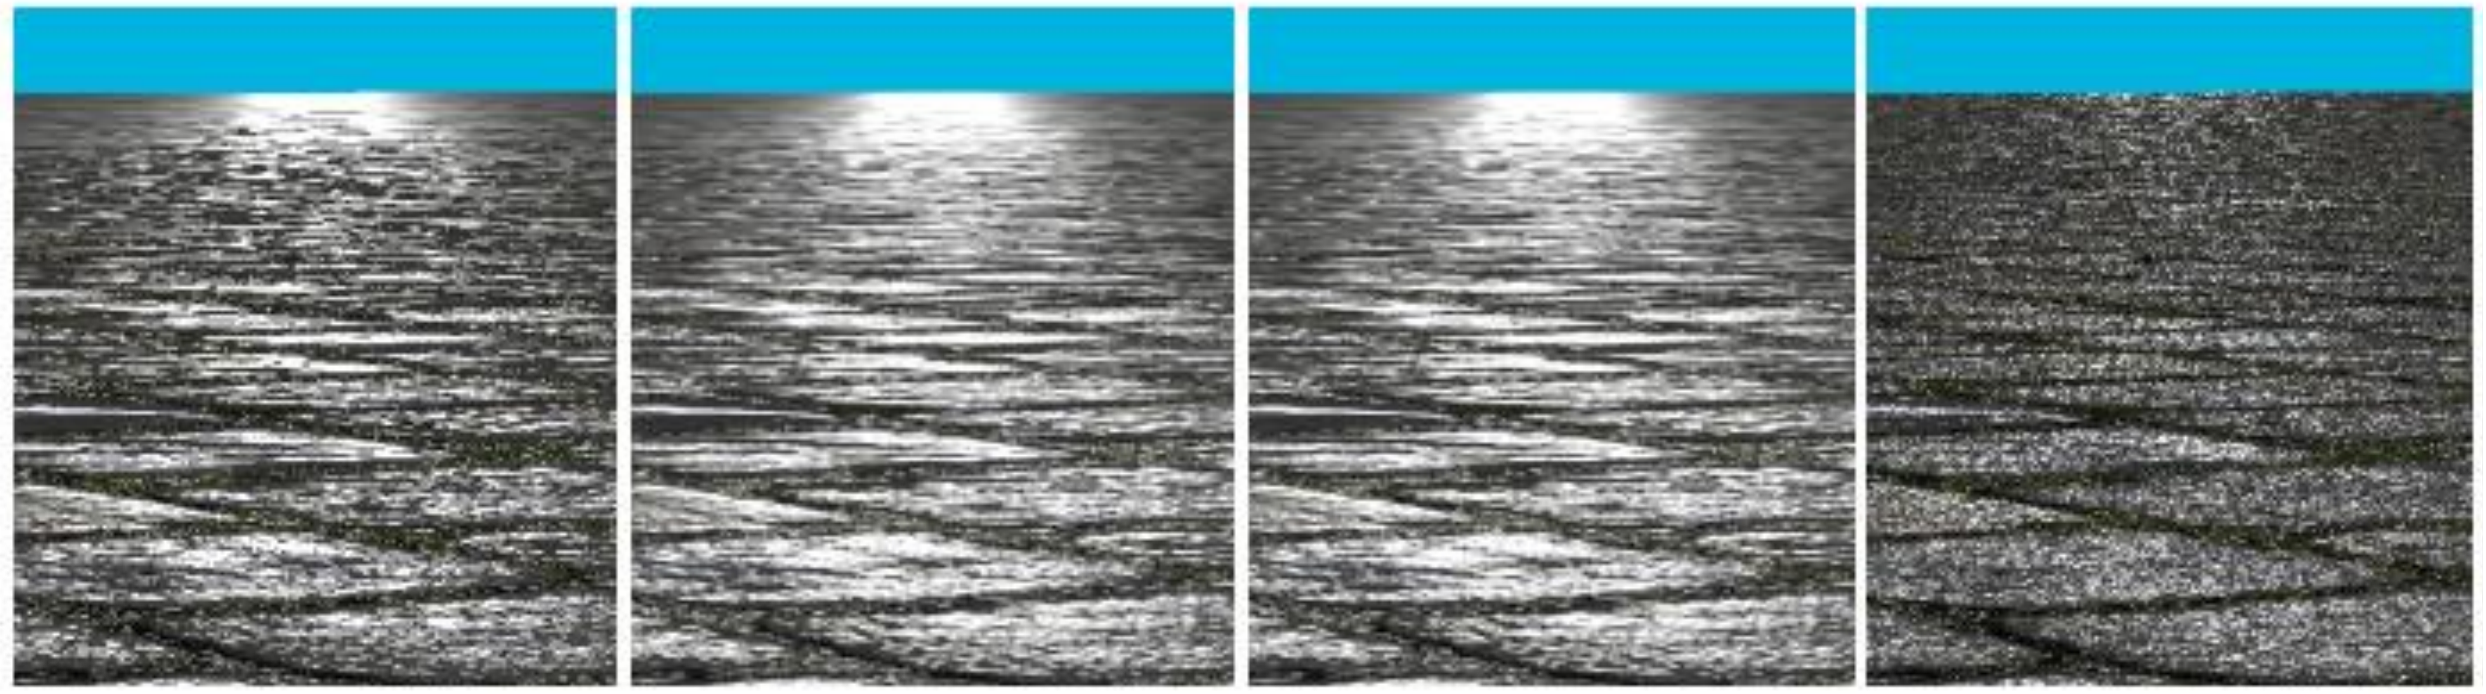

(a) Filter Before 1 spp | (b) Filter Before 1 spp + DLSS | (c) Filter Before 1024 spp | (d) Filter After 1 spp

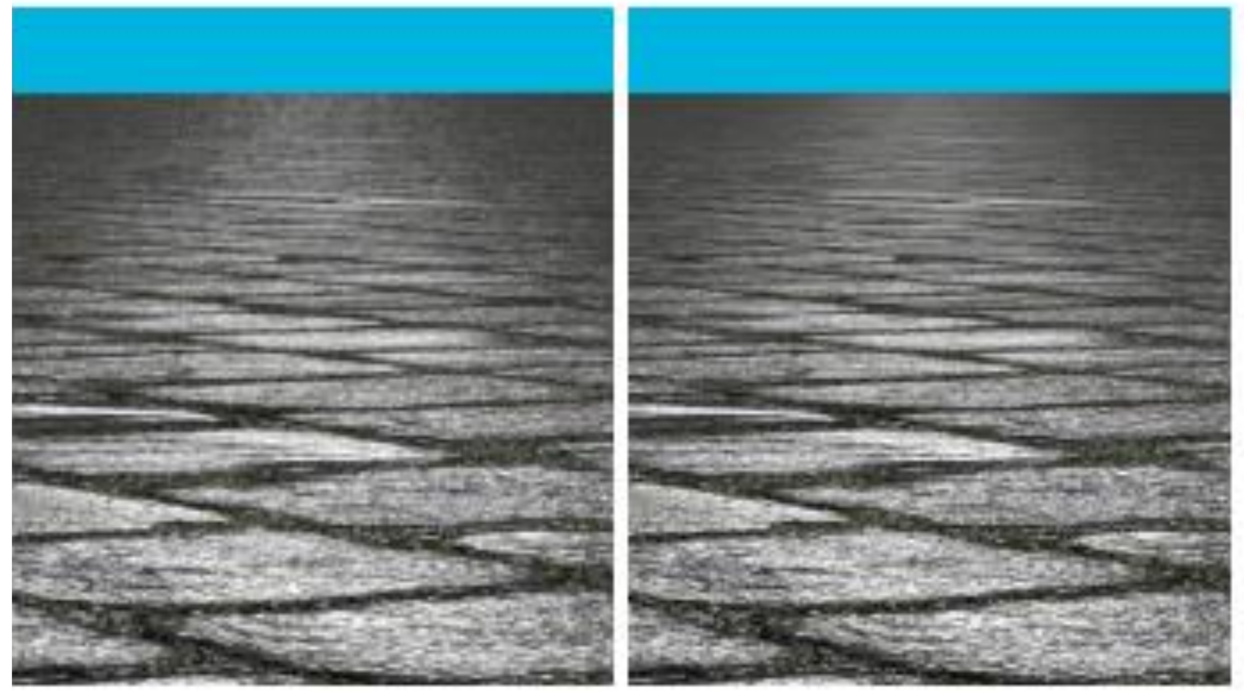

(g) Filter After LOD 0, 1 spp, DLSS | (h) Reference 1024 spp

Appearance preservation – real time

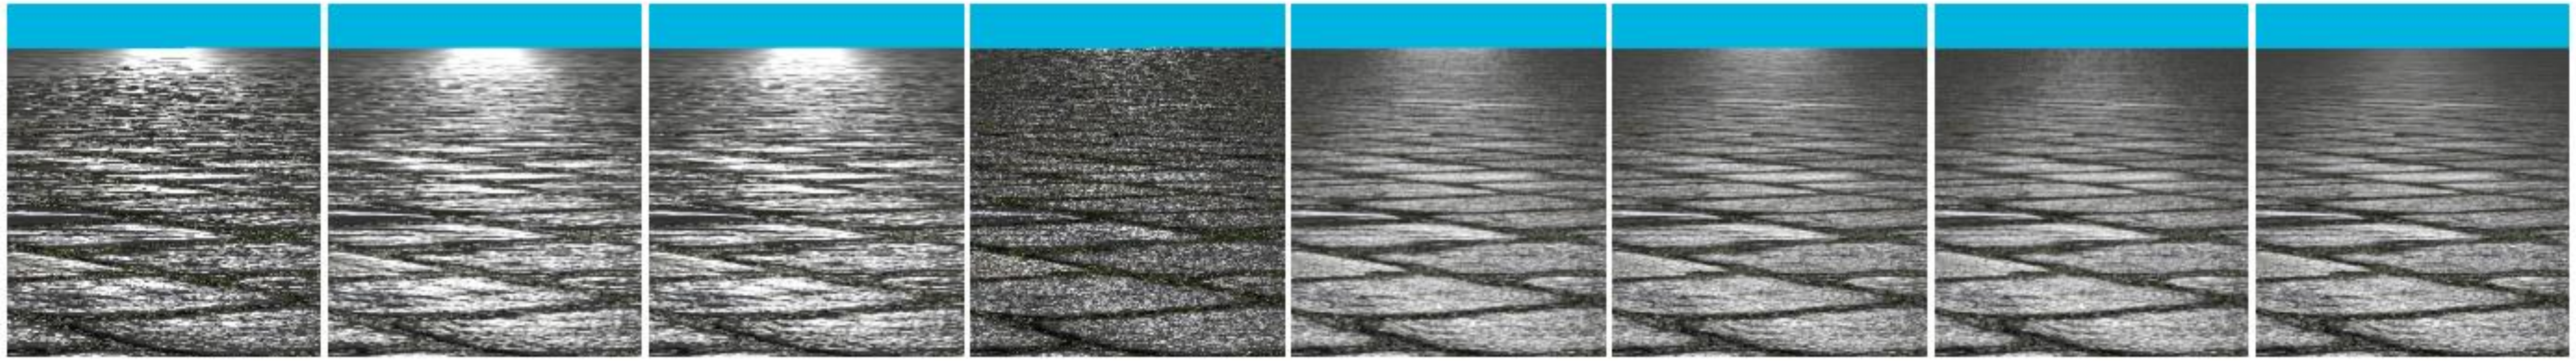

(a) Filter Before 1 spp | (b) Filter Before 1 spp + DLSS | (c) Filter Before 1024 spp | (d) Filter After 1 spp | (e) Filter After 1 spp + DLSS | (f) Filter After 1024 spp | (g) Filter After LOD 0, 1 spp, DLSS | (h) Reference 1024 spp

Hybrid: Use a mipmap, but higher resolution
(reduce cache trashing, make it easier for DLSS)

Appearance Preservation – offline, volumetric textures

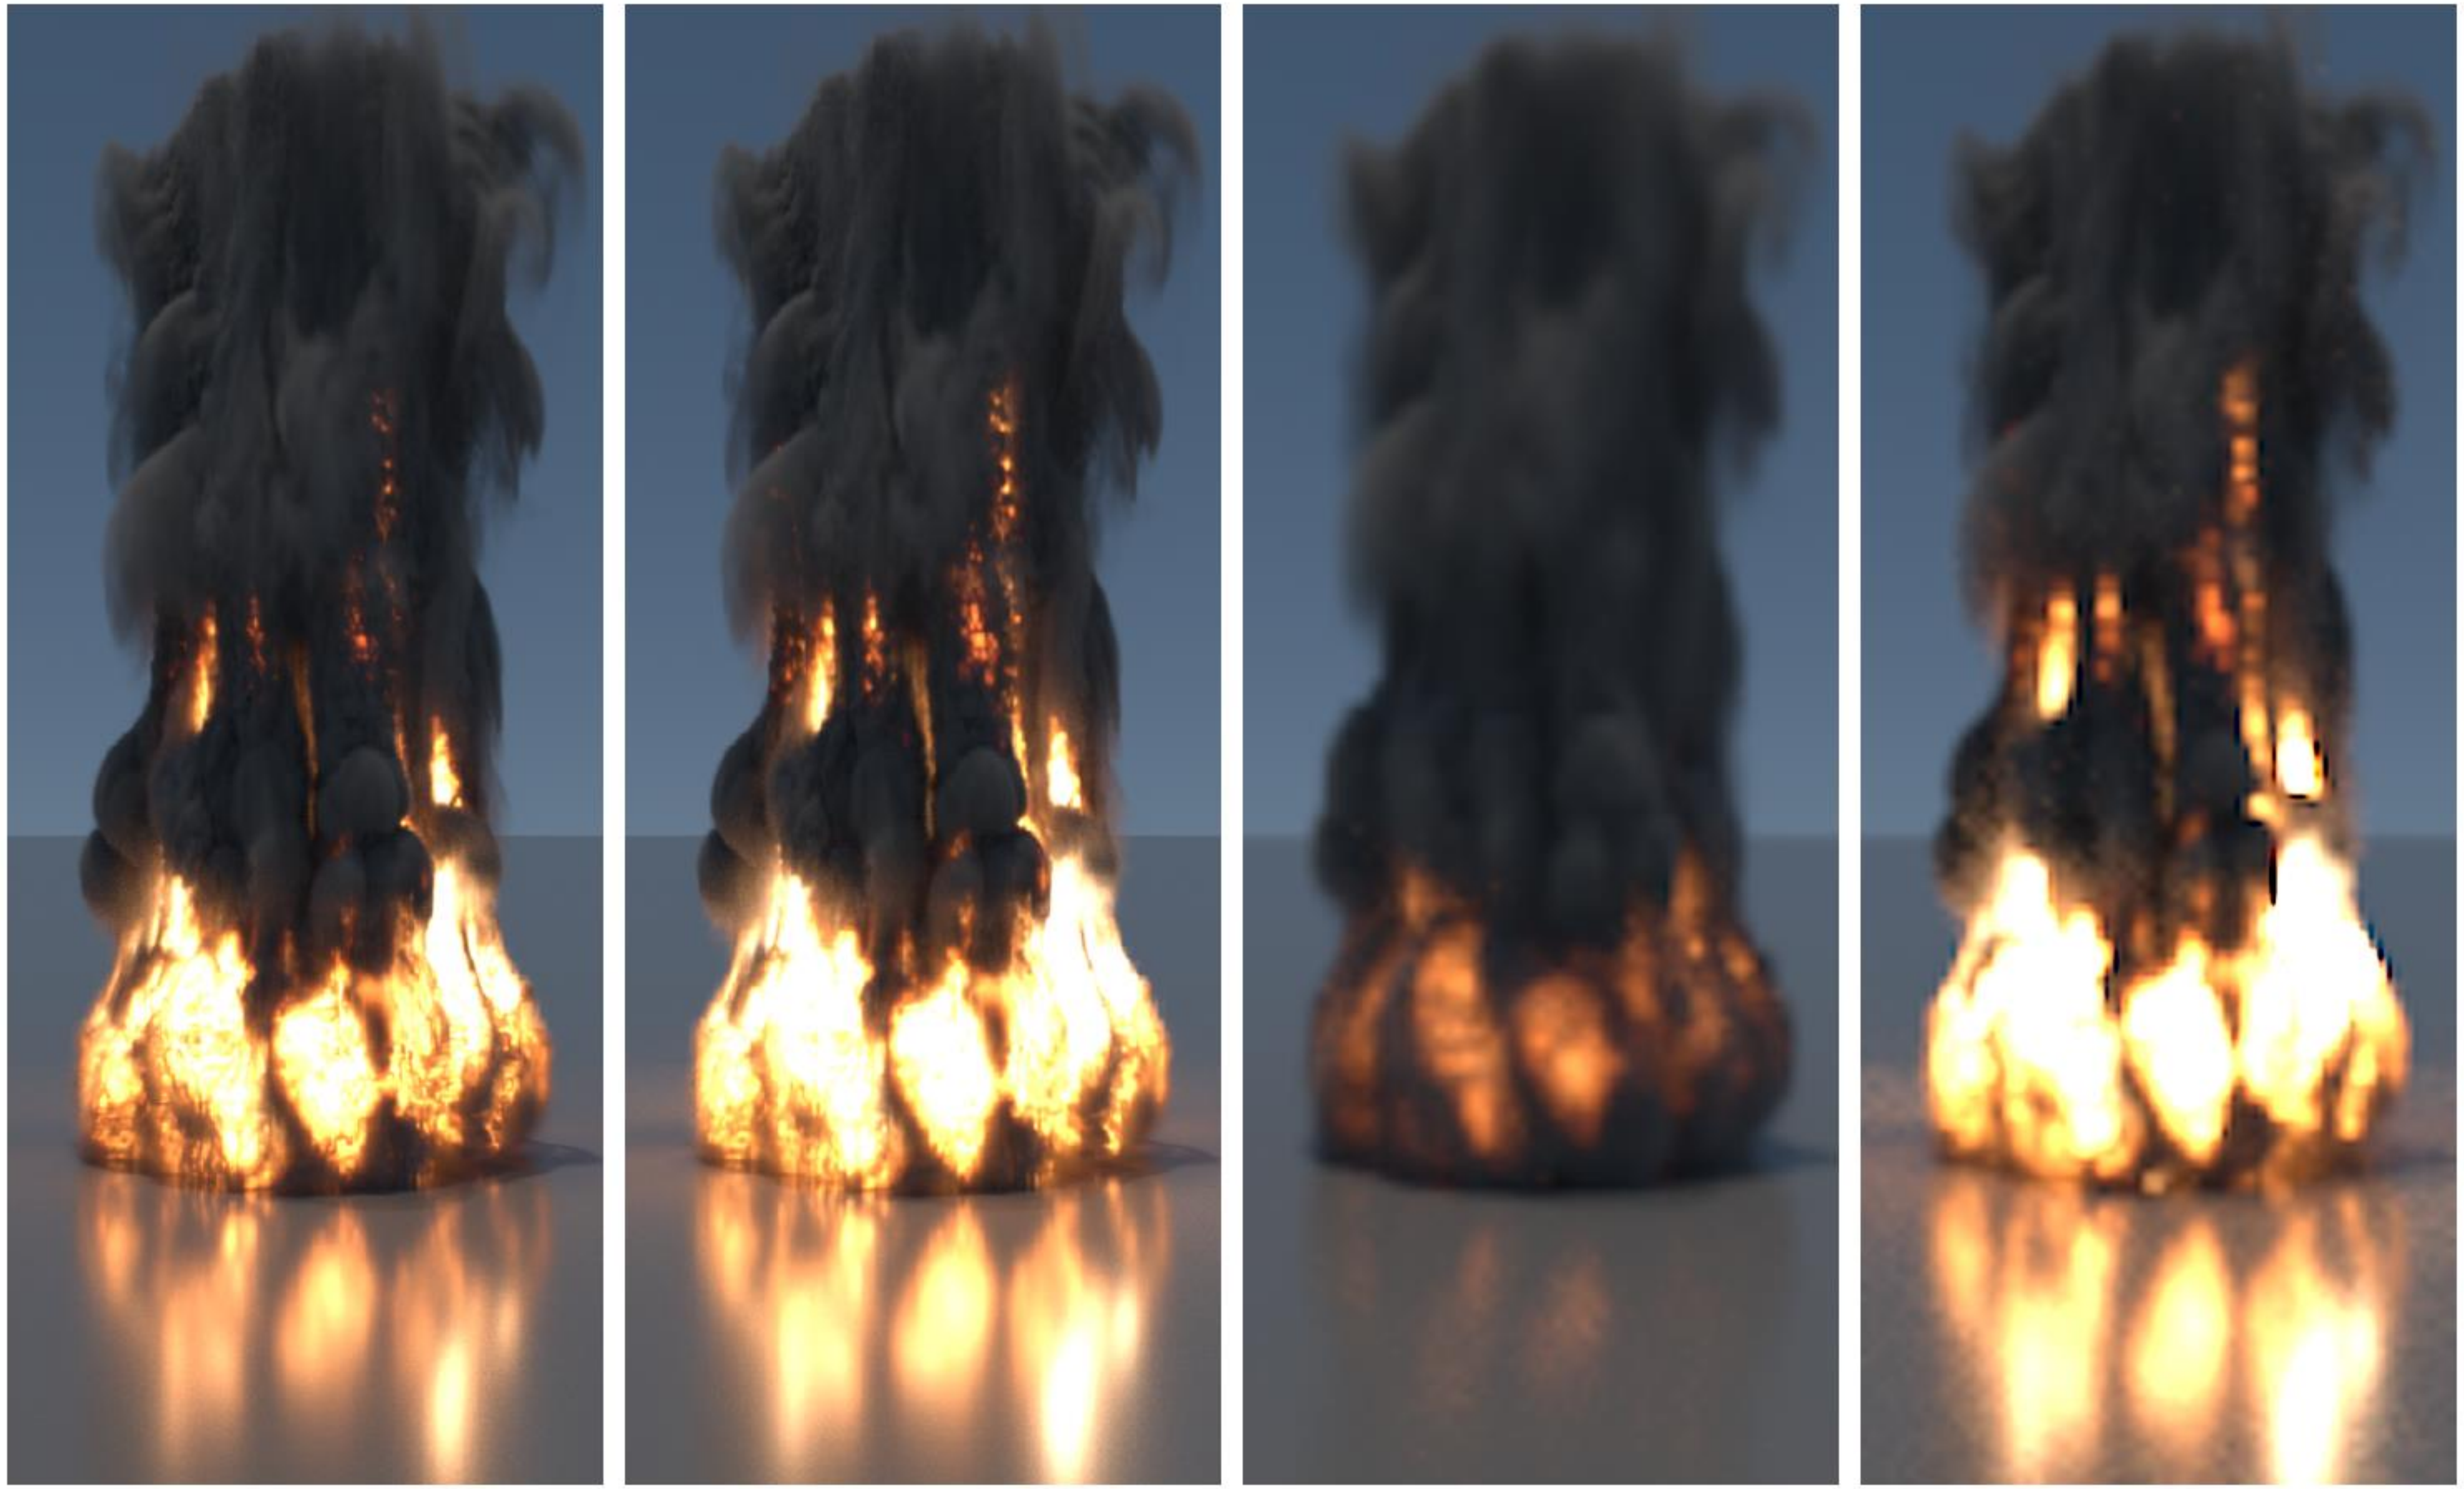

(a) Trilinear

(b) Stochastic trilinear

(c) Trilinear, MIP mapped

(d) Stochastic minification

Offline - Improved Image Quality & Performance
No additional noise!

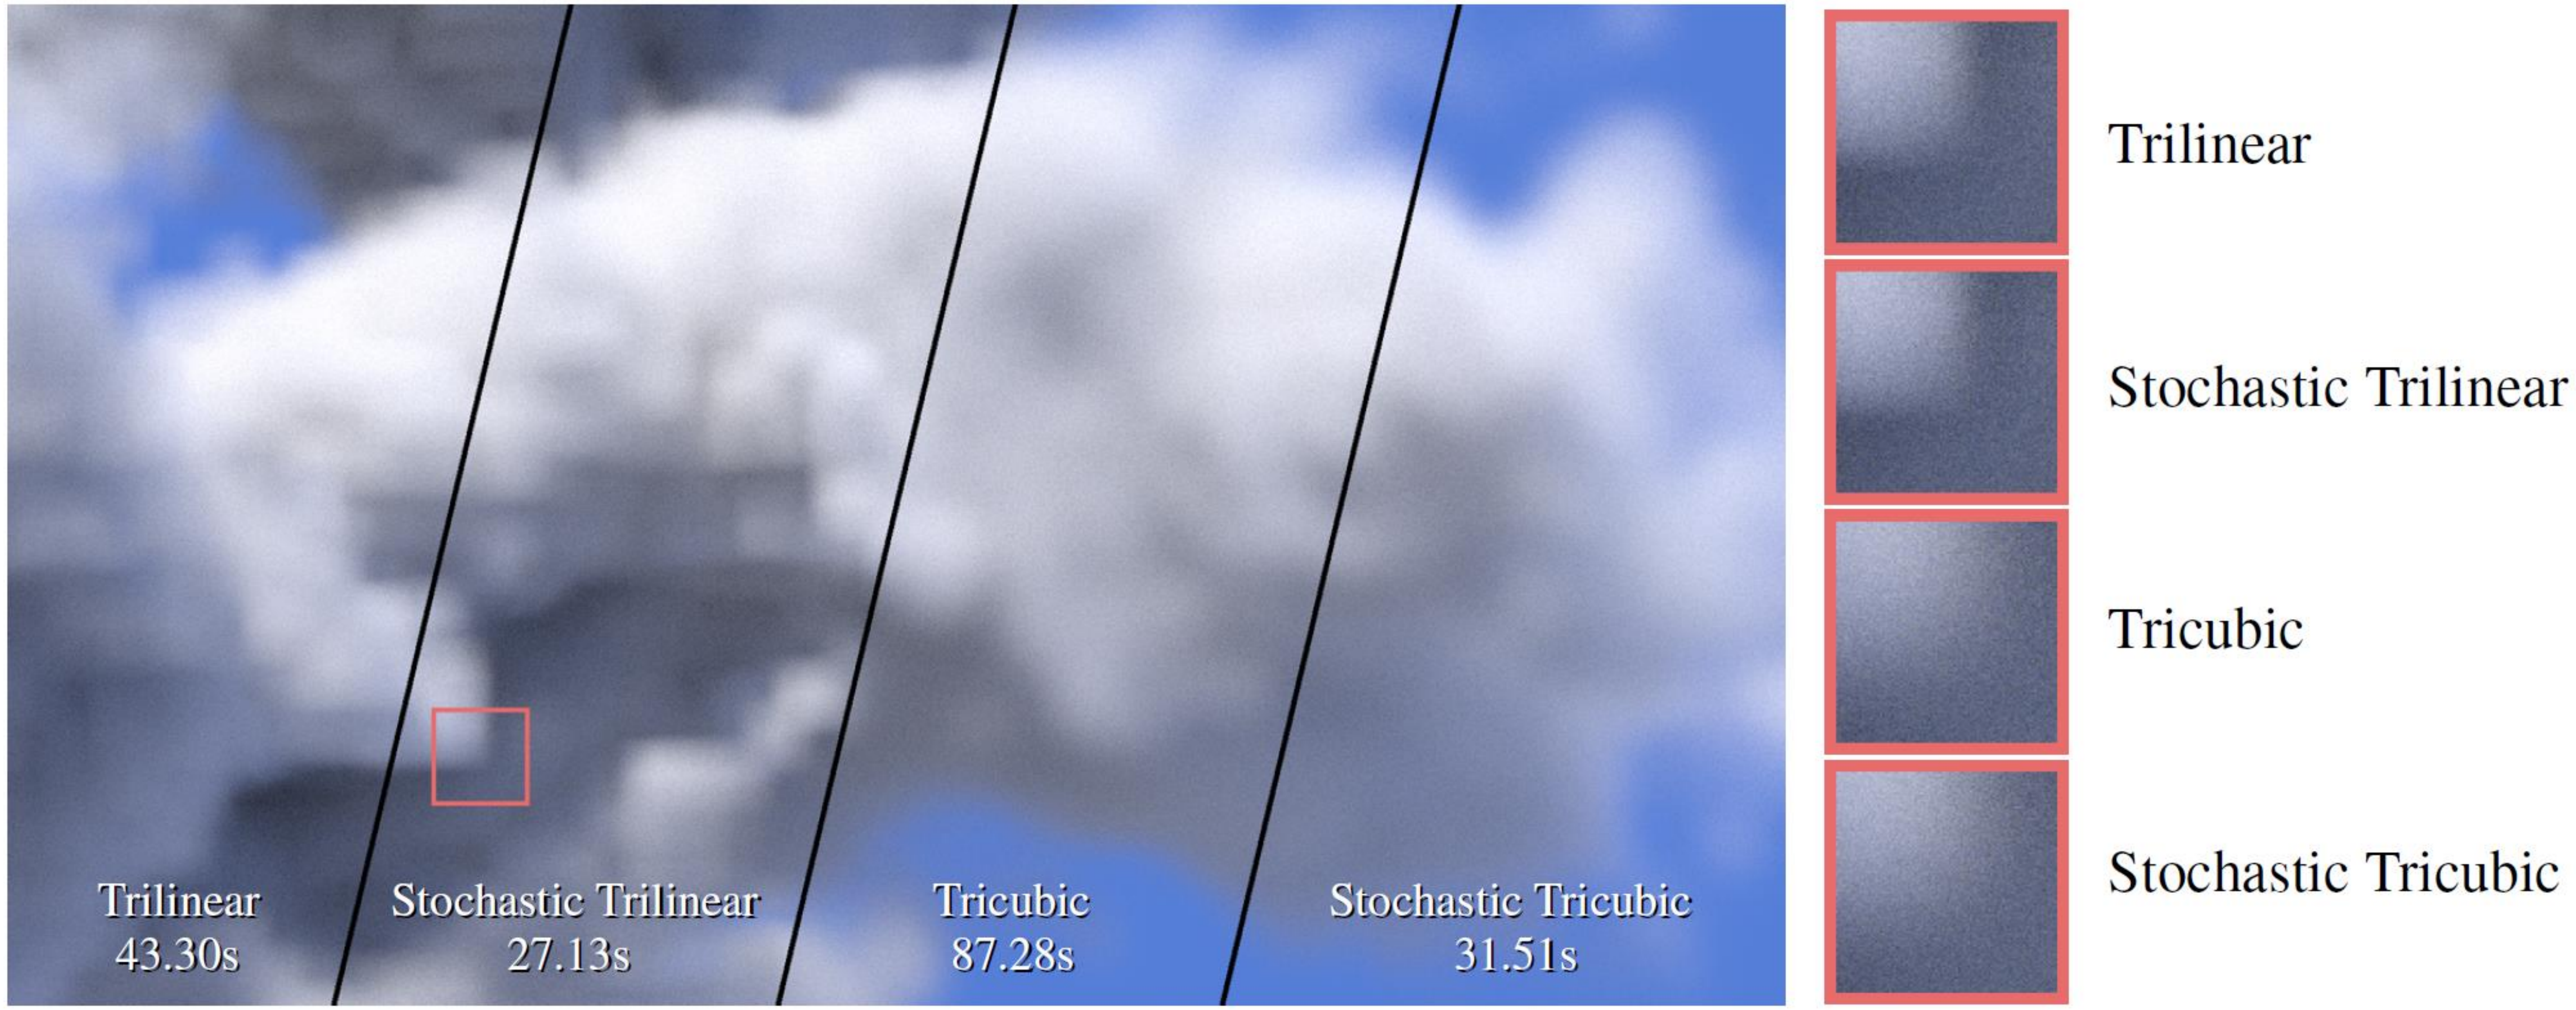

Figure 1: A section of the *Disney Cloud* scene rendered with path tracing. With this close-in viewpoint, trilinear filtering leads to blocky artifacts in the image. Tricubic filtering gives a much better result, though requires 64 voxel lookups into the *NanoVDB* representation. Stochastic filtering performs a single voxel lookup yet provides indistinguishable results, with overall rendering time speedups of 1.60× and 2.77× for the trilinear and tricubic filters. Times reported are for *pbrt-v4* running on an NVIDIA 4090 RTX GPU, rendering at 1080p with 256 samples per pixel.

Minification & Magnification
DLSS + STBN Temporal Stability Test

Stress Test: Real-Time Stochastic Filtering, high contrast, no mip-maps

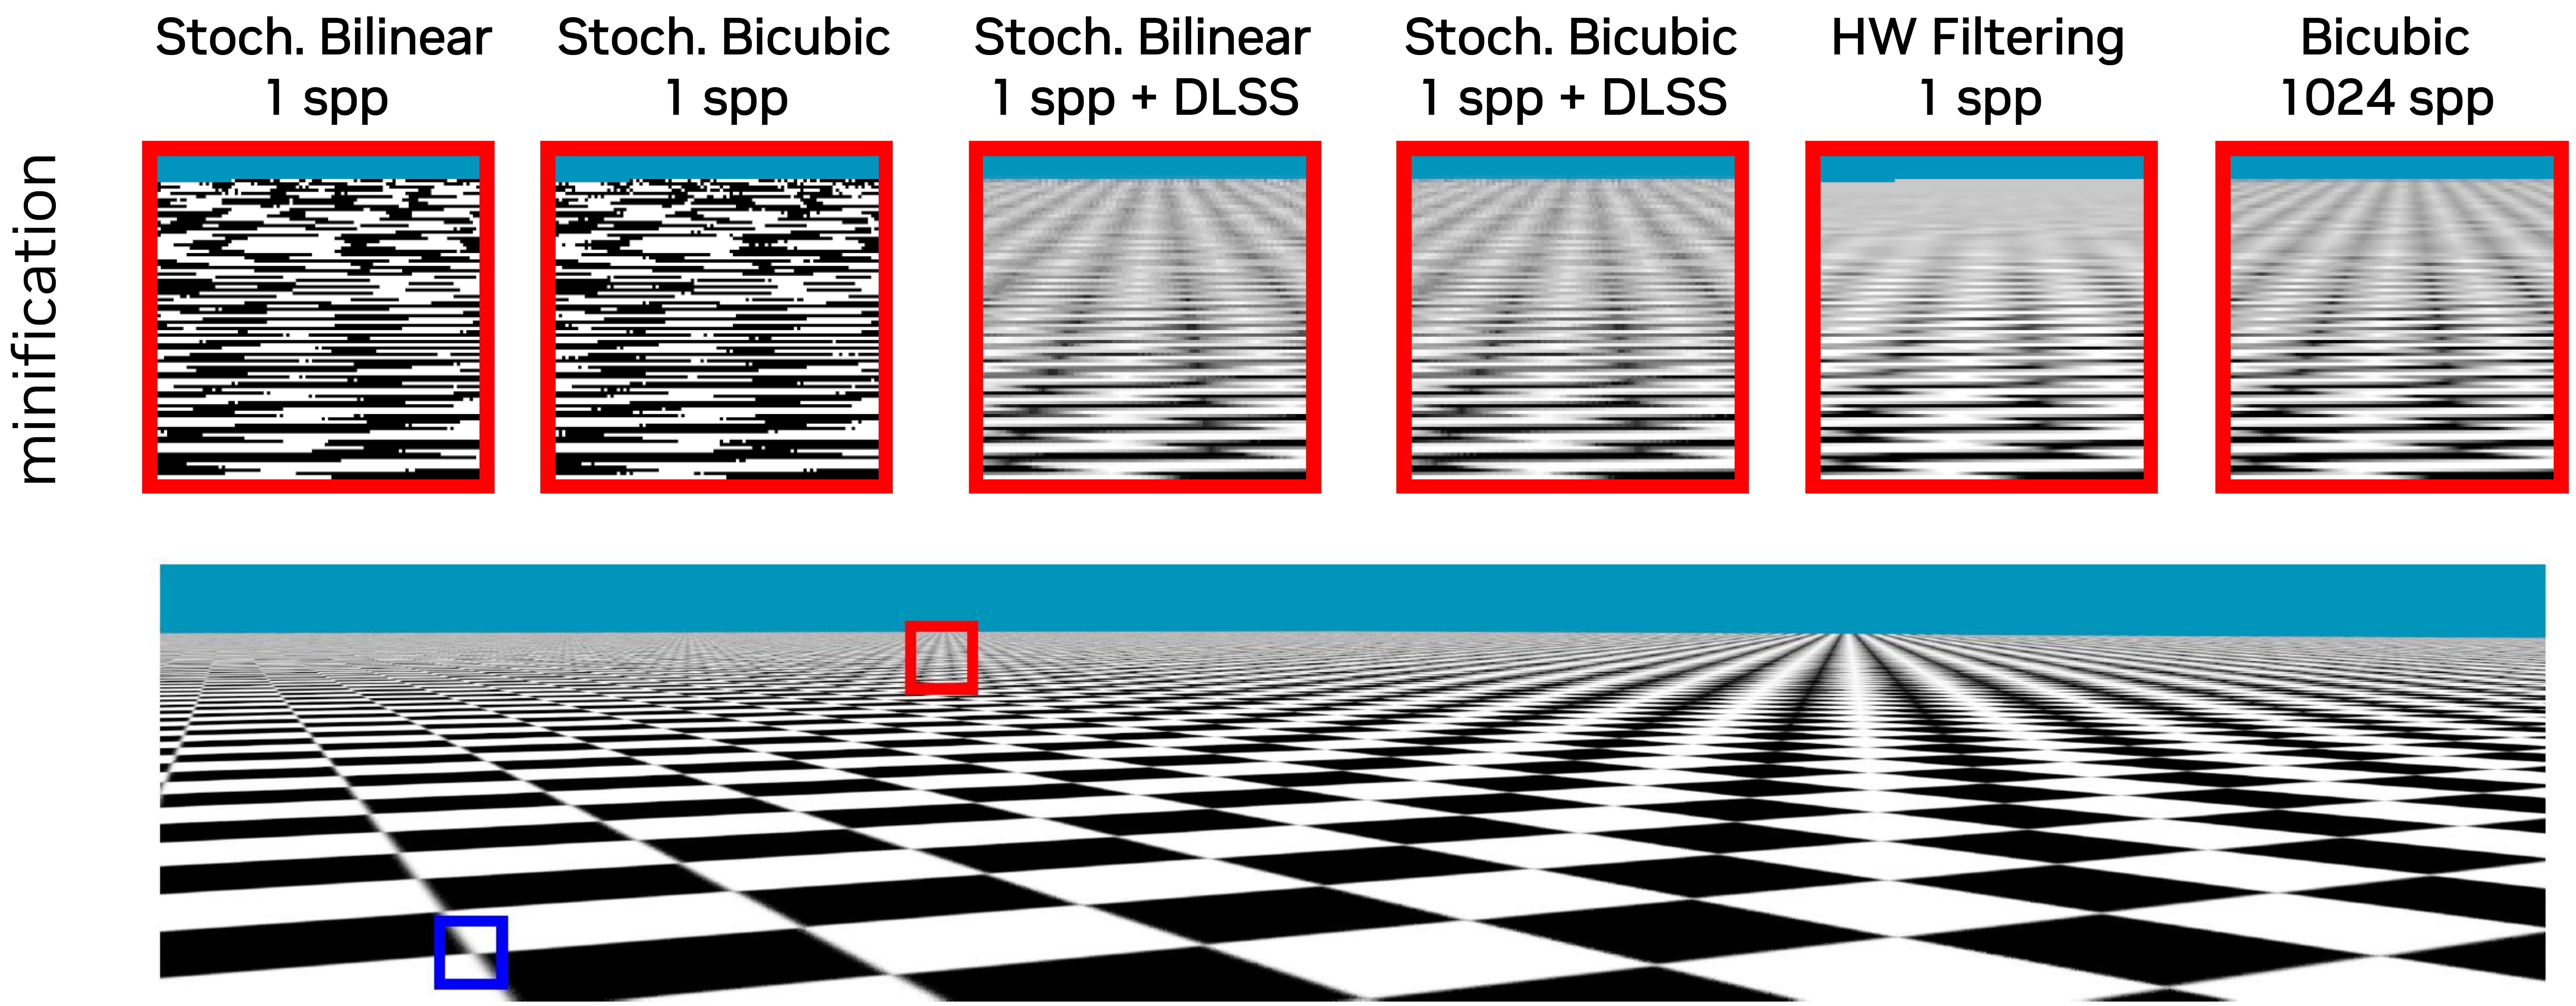

Stress Test: Real-Time Stochastic Filtering, high contrast, no mip-maps

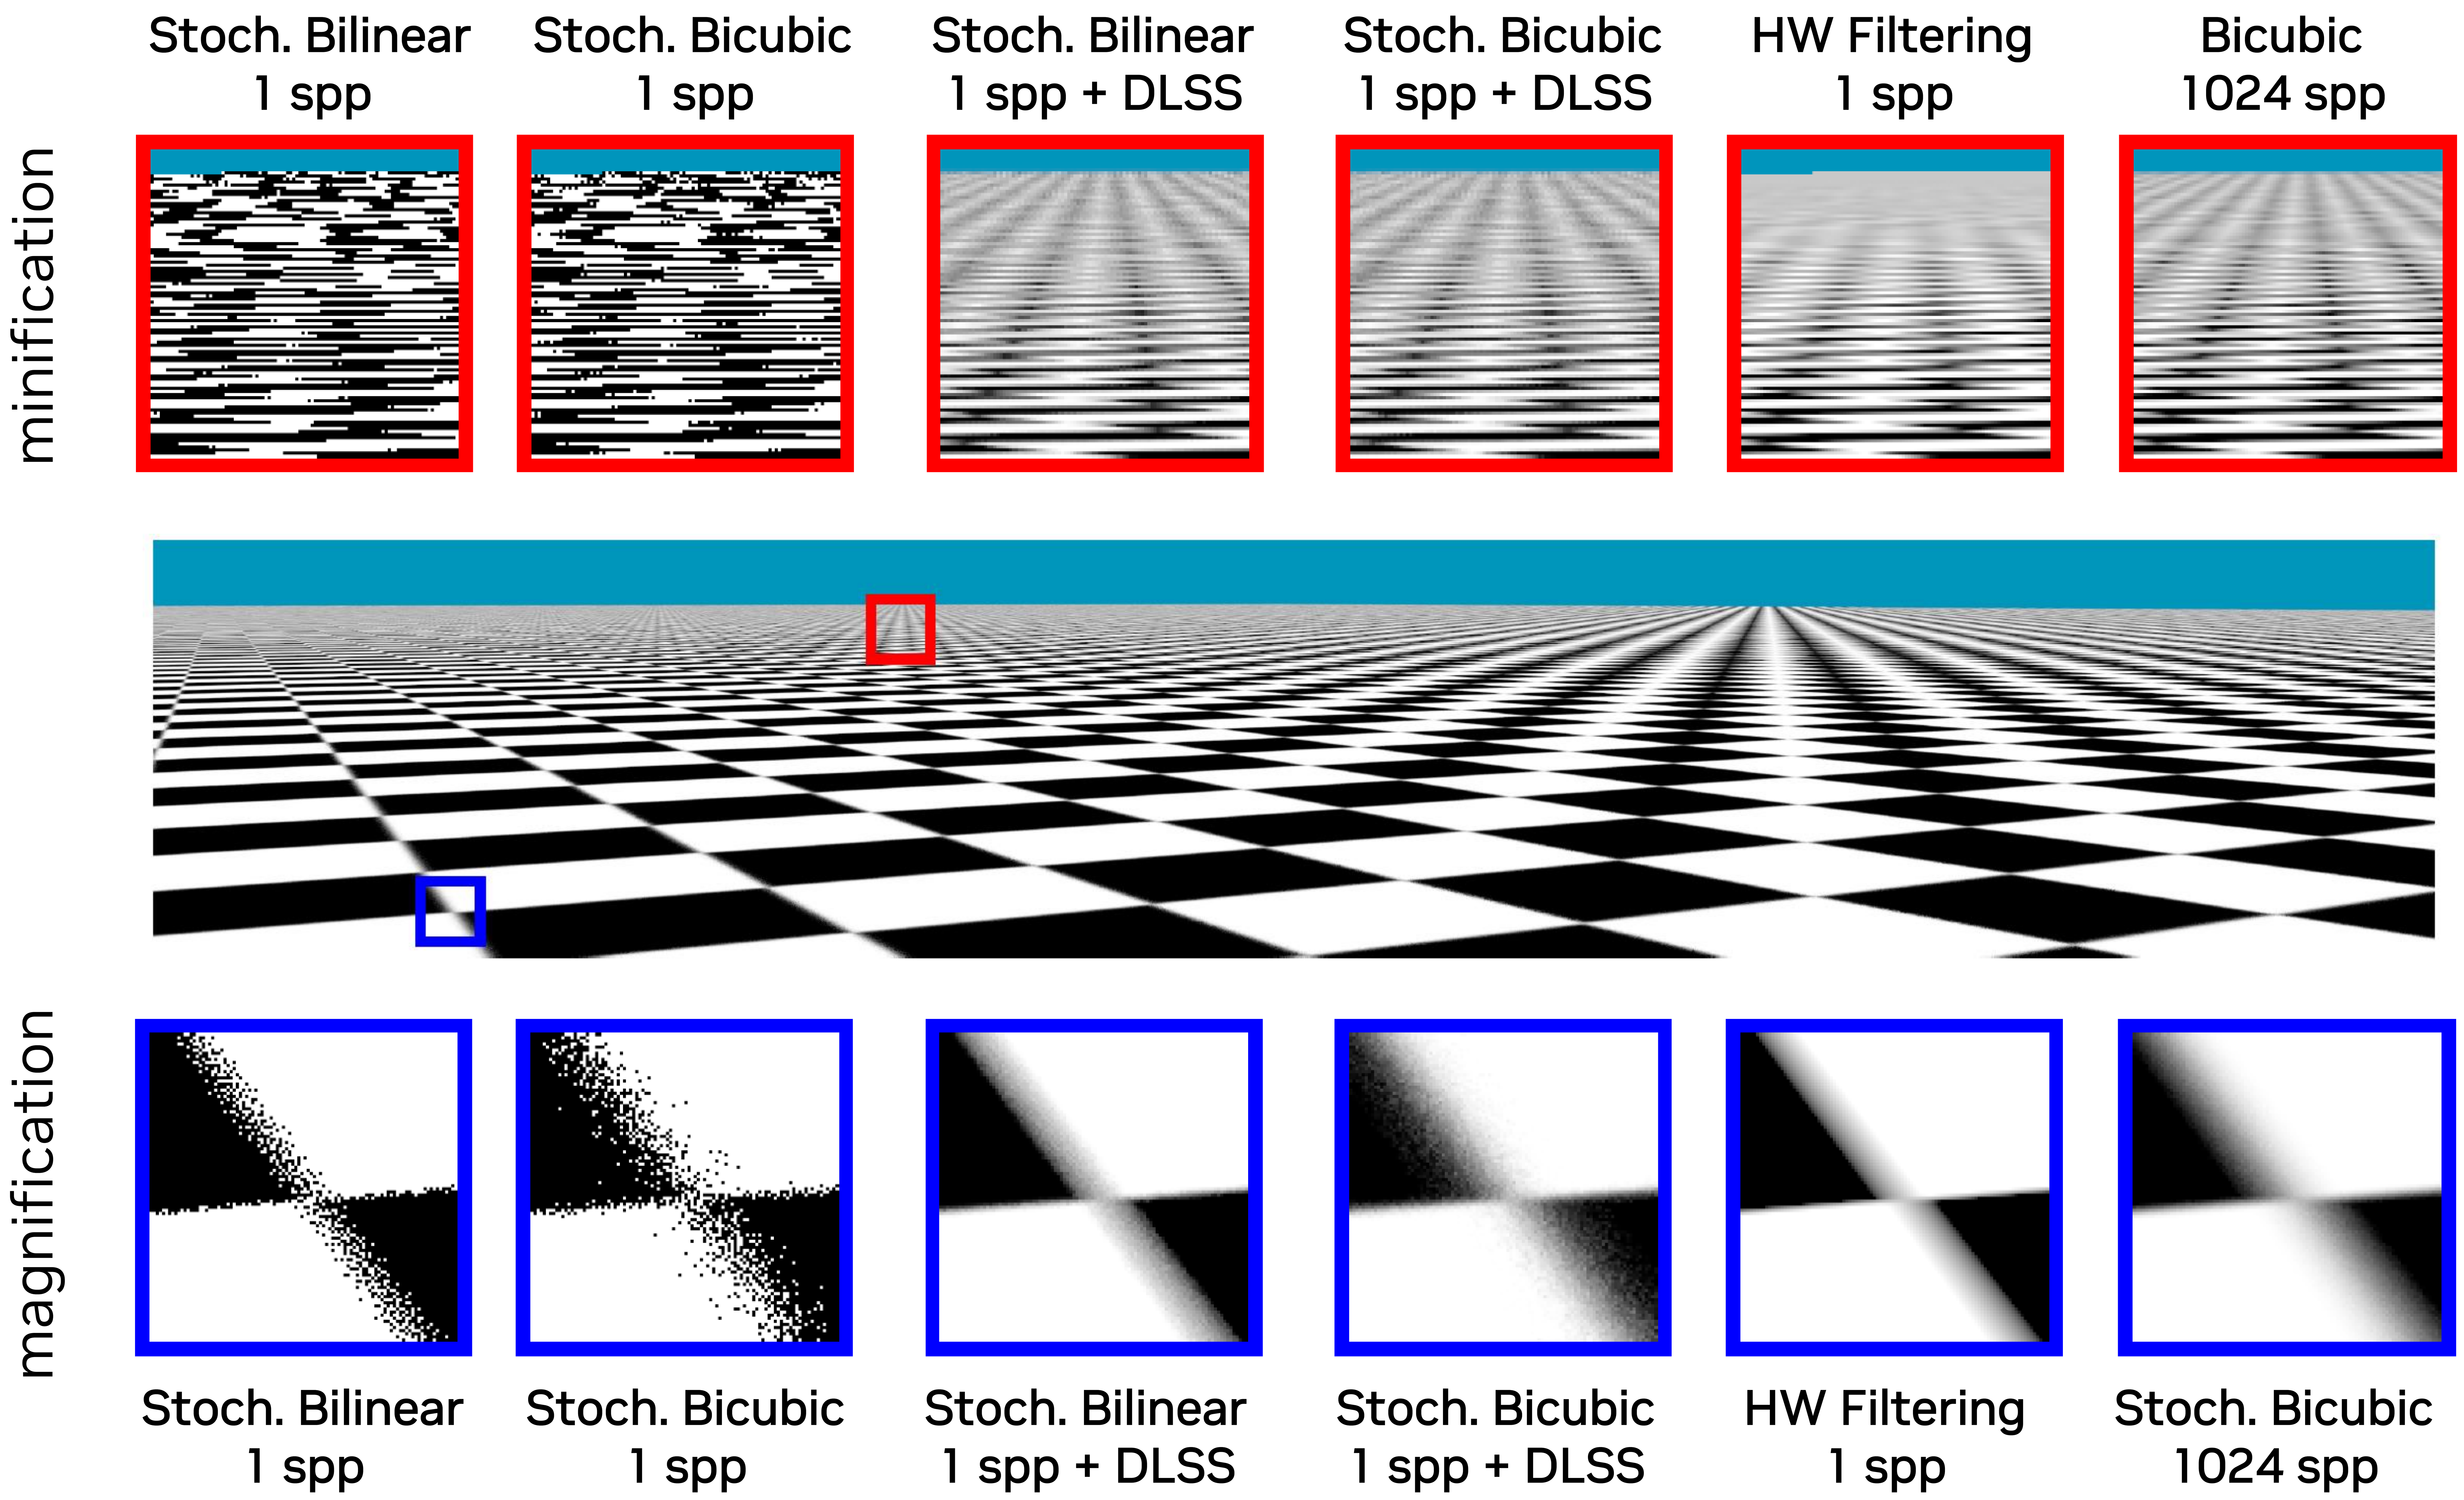

NVIDIA

Discussion

Noise

Even on the most noisy materials, does not flicker and is mild.
Future:

- Constant improvements to temporal reconstruction methods
- ML-based reconstruction can be trained
- STBN/LDS advances

Noise

Even on the most noisy materials, does not flicker and is mild.
Future:

- Constant improvements to temporal reconstruction methods
- ML-based reconstruction can be trained
- STBN/LDS advances

A general trend in graphics, two families of methods:

- Monte Carlo methods are general and can be unbiased, but initially too noisy and too slow
- Approximate, semi-analytical solutions used as a stopgap

Noise

Even on the most noisy materials, does not flicker and is mild.
Future:

- Constant improvements to temporal reconstruction methods
- ML-based reconstruction can be trained
- STBN/LDS advances

A general trend in graphics, two families of methods:

- Monte Carlo methods are general and can be unbiased, but initially too noisy and too slow
- Approximate, semi-analytical solutions used as a stopgap

Long term, Monte Carlo becomes practical and wins:

- TAA vs MSAA
- Shadowmap pre-filtering not very relevant today
- Path traced movies and even games!

Application – novel compression formats

NVIDIA

Application – novel compression formats

Not just NTC!

In the paper, we evaluate DCT

NeuralVDB -> octrees + NNs

Application – beyond filtering – material blending

Not just filtering... any interpolation!
- Includes material blending

Example: stochastic triplanar mapping
- Already practiced by game developers
- 3x faster
- Not just performance saving - unbiased!

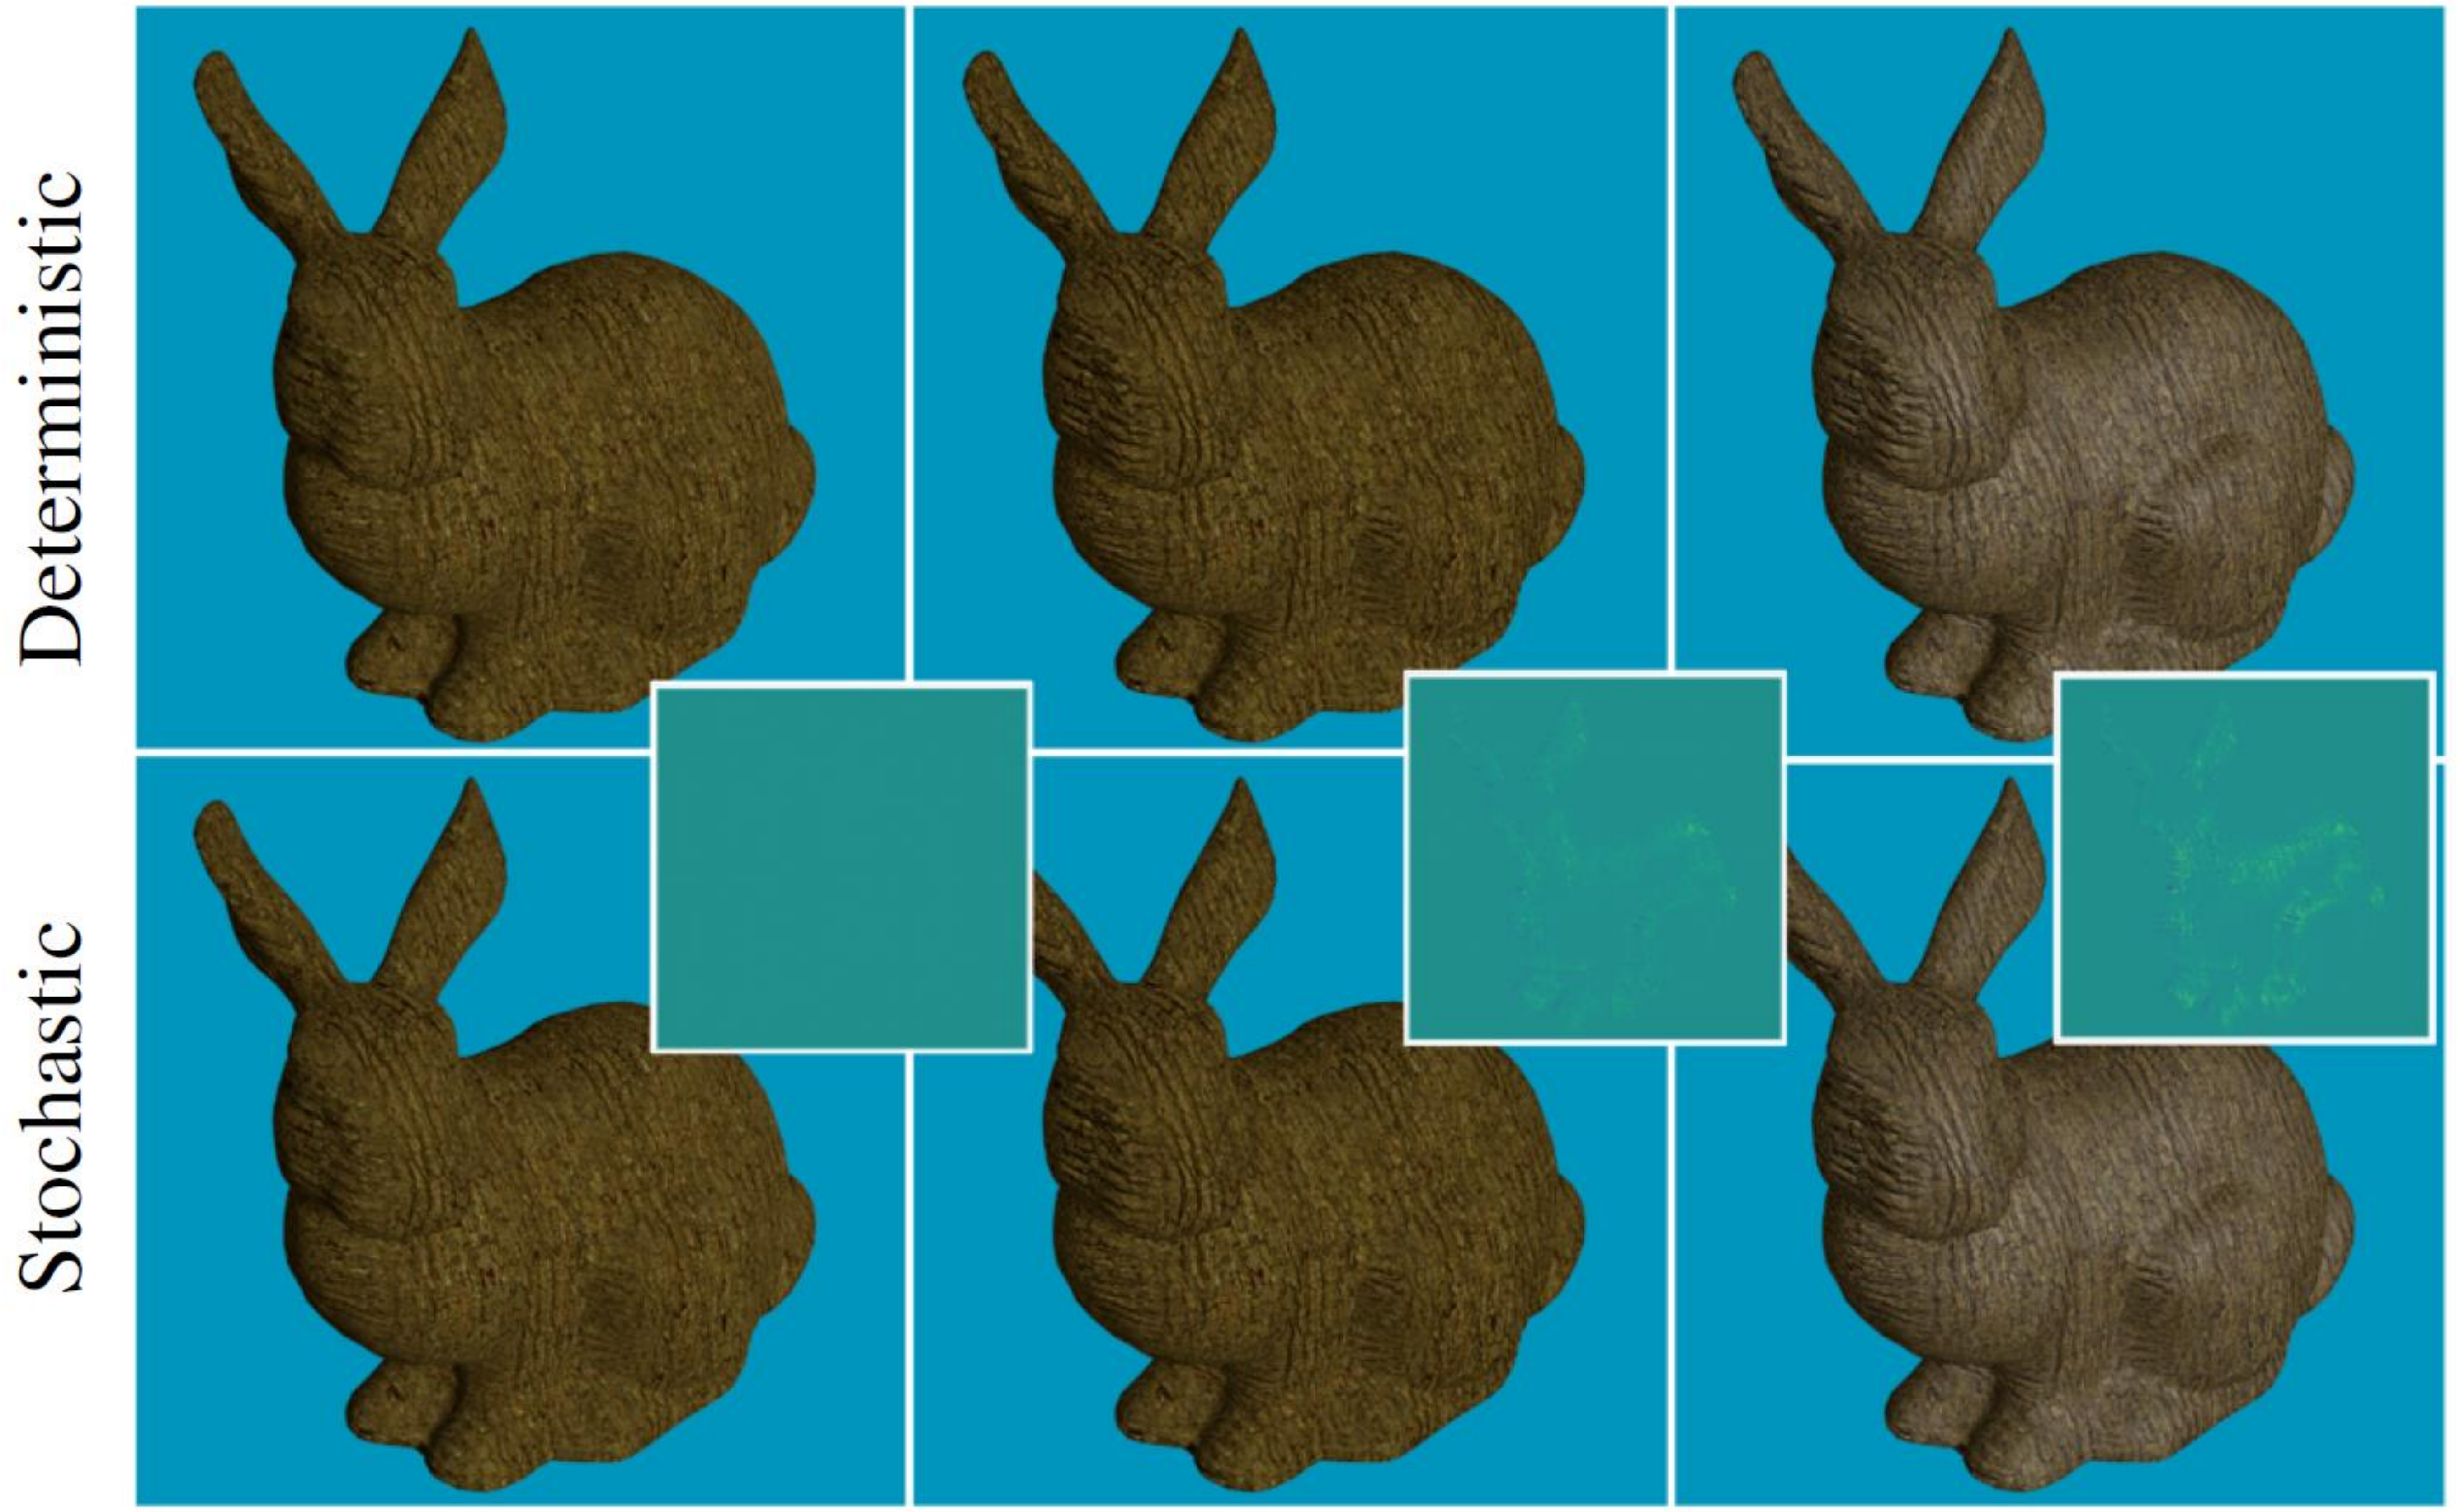

Figure 9: Full triplanar mapping (**top**) compared to its stochastic, single sample estimation (**bottom**). From left to right we present pure diffuse shading without normal mapping, diffuse shading with normal mapping, and full specular and diffuse lighting. Insets show error magnified 10×.

Application – better filters

Less aliased, sharper, smoother… up to you!

Bilinear

Mitchell

B-Spline Bicubic

Lanczos3

Recommendations

- **Minification**: Filtering After Shading is **always better**
- **Minification**: Offline rendering can remove mip-maps: rendering Monte Carlo noise dominates
- **Minification**: Real-time rendering: “hybrid” (performance, temporal stability but some bias remains)

Recommendations

- **Minification**: Filtering After Shading is **always better**
- **Minification**: Offline rendering can remove mip-maps: rendering Monte Carlo noise dominates
- **Minification**: Real-time rendering: “hybrid” (performance, temporal stability but some bias remains)

- **Magnification**: Filtering After Shading is **unbiased and removes errors**
- **Magnification**: Filtering After Shading **simplifies logic** (alpha, metalness, texture padding)
- **Magnification**: Filtering After Shading **allows for better filters and new texture representations**

Recommendations

- **Minification**: Filtering After Shading is **always better**
- **Minification**: Offline rendering can remove mip-maps: rendering Monte Carlo noise dominates
- **Minification**: Real-time rendering: “hybrid” (performance, temporal stability but some bias remains)

- **Magnification**: Filtering After Shading is **unbiased and removes errors**
- **Magnification**: Filtering After Shading **simplifies logic** (alpha, metalness, texture padding)
- **Magnification**: Filtering After Shading **allows for better filters and new texture representations**
- **Magnification**: Filtering After Shading **can introduce aliasing**
- **Magnification**: It depends! Decide based on use-case, content type, maximum magnification

Recommendations

Magnification: If you can get 16x better compression by using novel compression format and STF...

Having 16x more real texels is better than relying on interpolation!

Recommendations

Magnification: If you can get 16x better compression by using novel compression format and STF…
Having 16x more real texels is better than relying on interpolation!

- You don’t have to go “all in”, we recommend a pragmatic approach:
    - There are trade-offs and cases where one is preferred over the other
    - Don’t stochastically sample something that relies on interpolation (e.g., SDF fonts)
    - Use STF/non-STF/different filters on different assets – only shader code changes!

Conclusions

- Our proposal of “**filtering after shading**” might seem radical…
- We simply formalize decades of the different film industry and gamedev practices
    - Filtering after shading is unbiased and better for appearance preservation
    - We need to change the way we teach filtering and *blending*

Conclusions

- Our proposal of “filtering after shading” might seem radical…
- We simply formalize decades of the different film industry and gamedev practices
  - Filtering after shading is unbiased and better for appearance preservation
  - We need to change the way we teach filtering and *blending*

- **Stochastic texture filtering** present for ~40y in literature in various one-off flavors
  - We explain the prior approaches and generalize them
  - We propose two families of techniques with different trade-offs

Conclusions

- Our proposal of “filtering after shading” might seem radical…
- We simply formalize decades of the different film industry and gamedev practices
    - Filtering after shading is unbiased and better for appearance preservation
    - We need to change the way we teach filtering and *blending*

- **Stochastic texture filtering** present for ~40y in literature in various one-off flavors
    - We explain the prior approaches and generalize them
    - We propose two families of techniques with different trade-offs
    - We expand those to more filters, including negative lobe filters
    - Source code of efficient implementations – drop-in, zero integration cost!

Summary

Filtering After Shading by Stochastic Texture Filtering is a valuable tool:

- Remove workarounds and simplify code
- Enables efficient filtering of novel compression and storage formats
- Efficient and better filters
- Beyond textures: optimize and stochastically sample complex shader graphs

It is practical today! Just try it out. ☺

NVIDIA

Acknowledgments

Tomas Akenine-Möller
Homam Bahnassi
John Burgess
Johannes Deligiannis
Nathan Hoobler
Markus Kettunen
Aaron Lefohn
Peter Morley
Alexey Panteleev
Rajeev Penmatsa
Karthik Vaidyanathan
...and more – thank you!

Thank you for listening!

https://research.nvidia.com/publication/2024-05_filtering-after-shading-stochastic-texture-filtering

Bonus Slides

Enable Custom Texture Compression/Storage Algorithms

Minification vs Magnification jitter

Bonus: unexpected consequence

- Something that bothered me for many years…
- We always recommend decoding to linear before generating mip-maps (minification)…
- But why upsampling/sharpening looks way better applied in sRGB/gamma space?
- **Gamma conversion in either direction – introduces aliasing!**
- Doing/undoing gamma correction: Alias -> upsample -> Alias

Upsampling in sRGB

Upsampling in linear

Figure credit*: A Fresh Look at Generalized Sampling,* Diego Nehab and Hugues Hoppe

Sampling Texture Filters – Negative Lobes

- Image Processing uses almost exclusively negative lobe filters
- Approximations of a “perfect” interpolation filter
- Sharp, anti-aliased

B-Spline Bicubic

Mitchell

Lanczos3

Sampling Texture Filters – Negative Lobes

- Image Processing uses almost exclusively negative lobe filters
- Approximations of a “perfect” interpolation filter
- Sharp, anti-aliased
- Examples: Sinc, Lanczos, Mitchell…

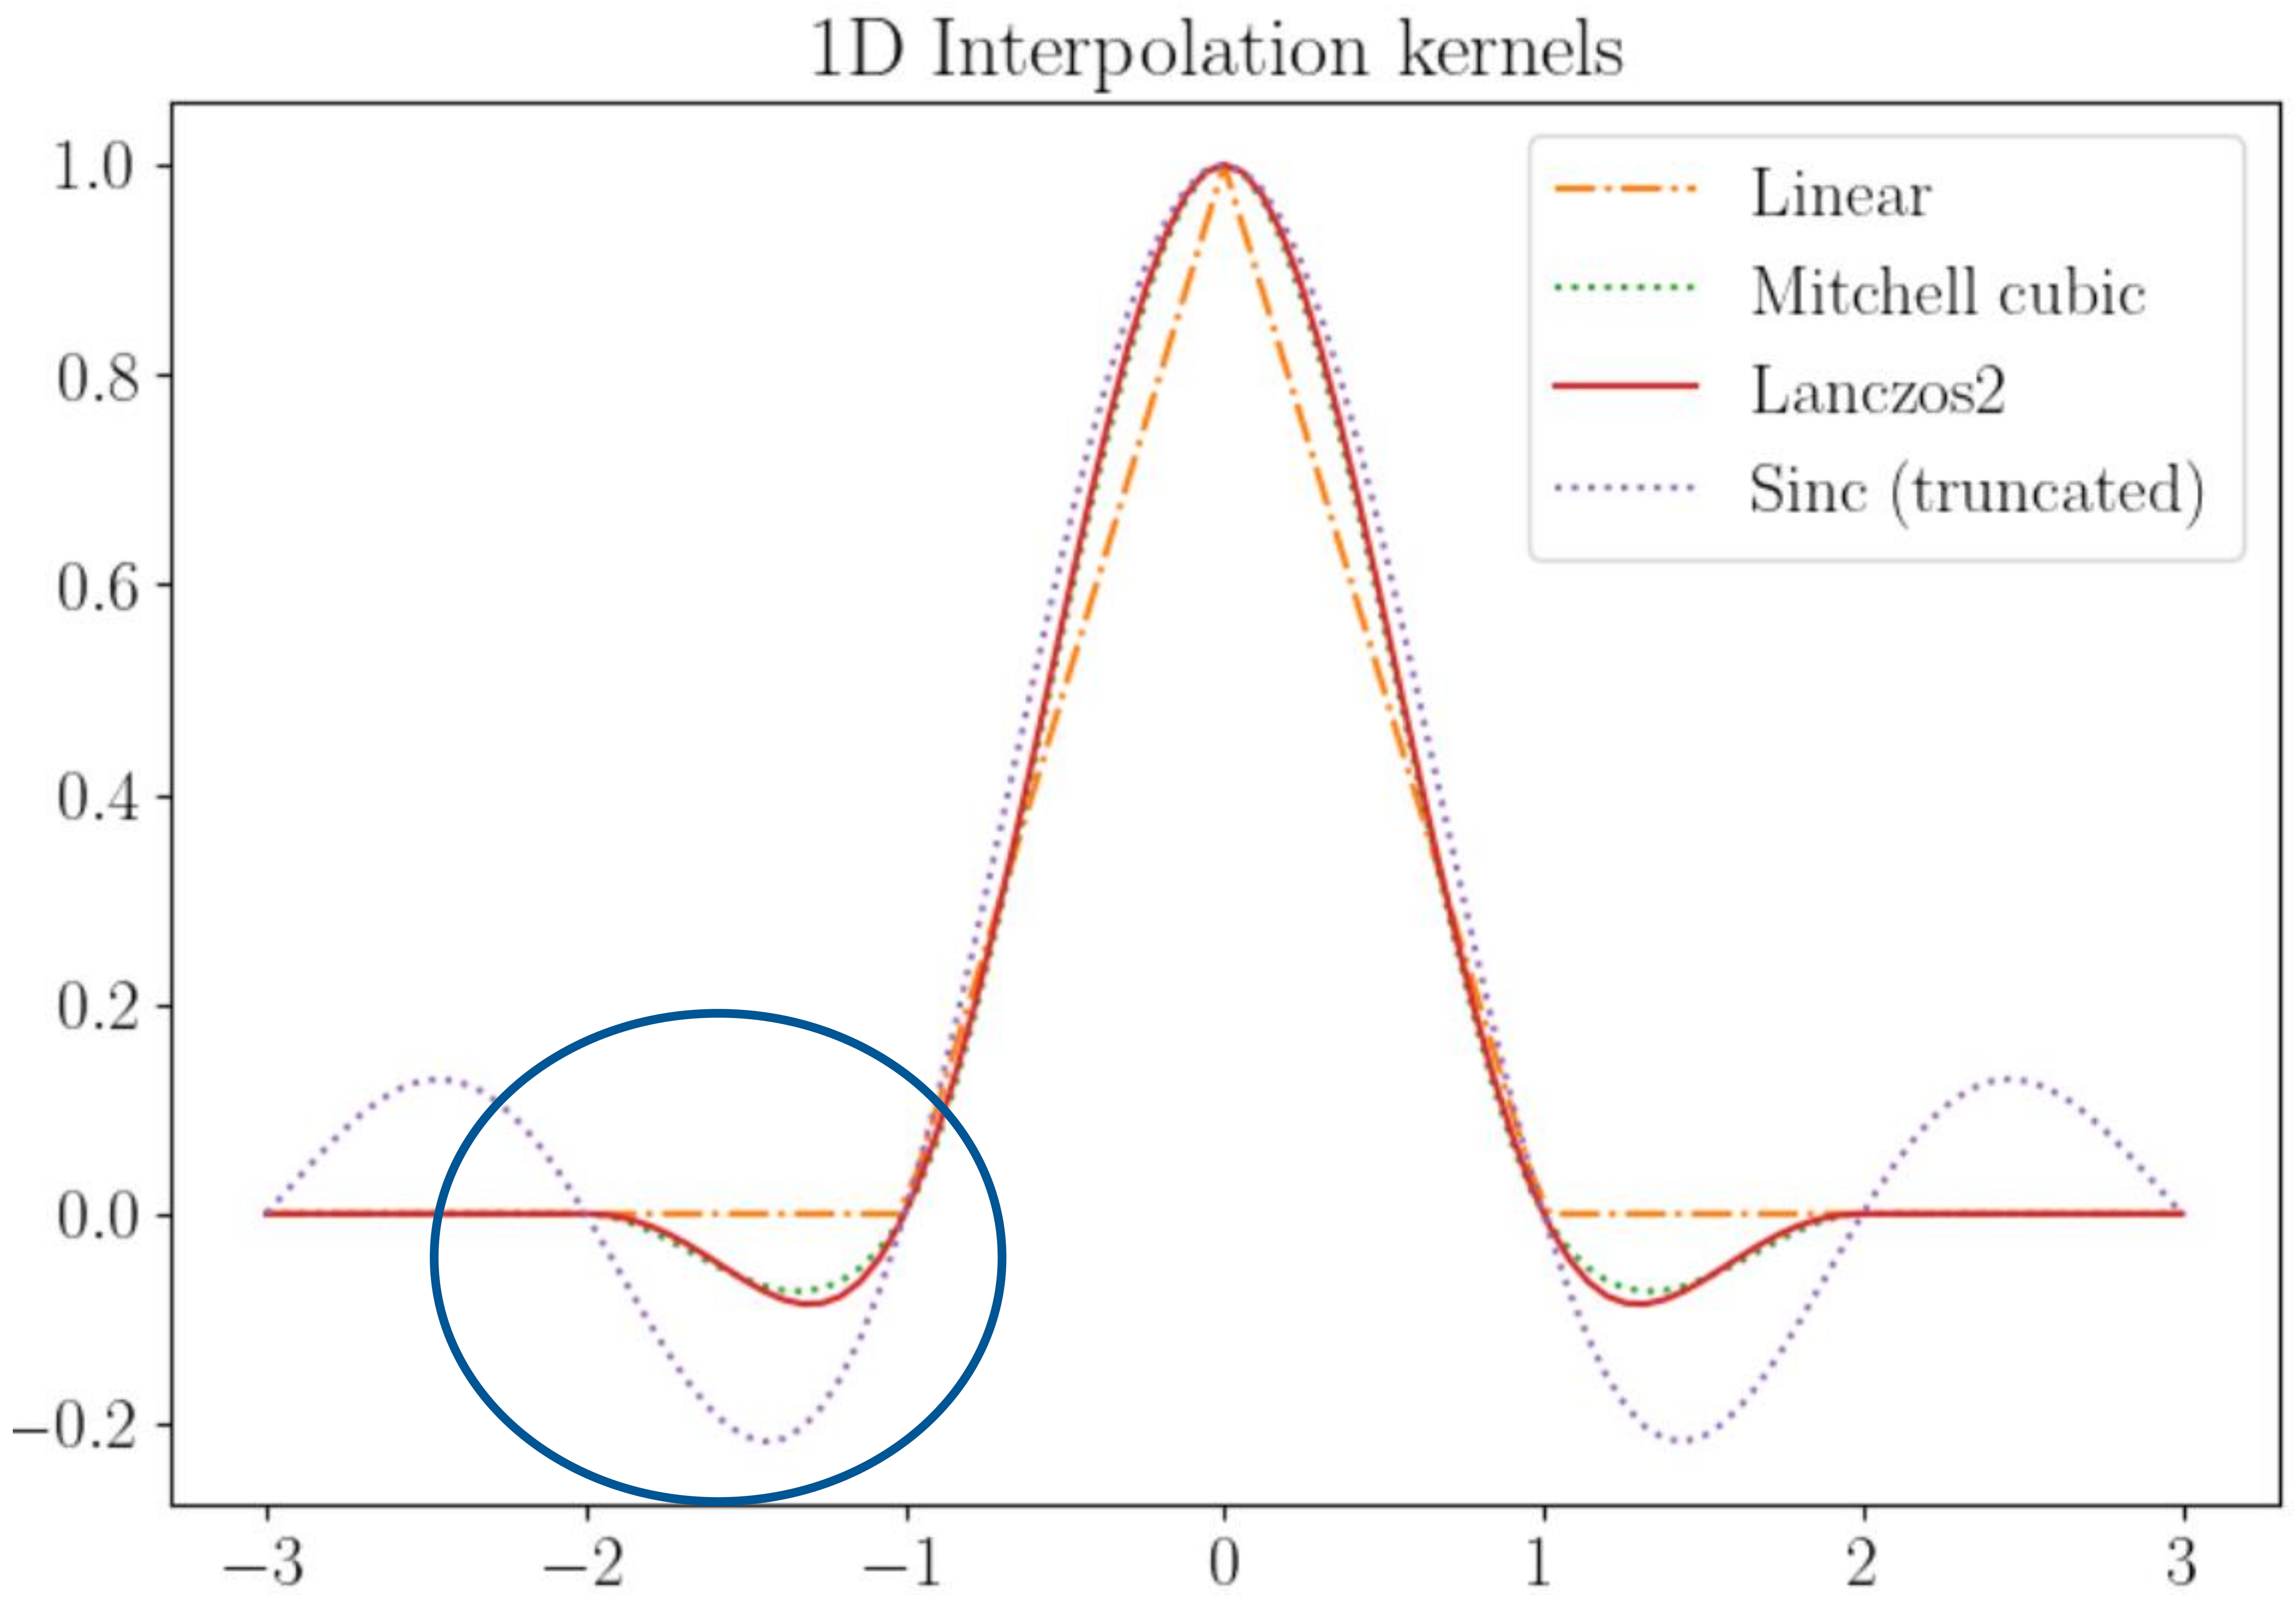

Sampling Texture Filters – Negative Lobes

- Sample proportionally to abs(f) -> works, but…
- Generates negative values
- Very high variance and noise

Sampling Texture Filters – Negative Lobes

- Sample proportionally to abs(f) -> works, but…
- Generates negative values
- Very high variance and noise

Solution – **positivization**

- Importance sample the positive and negative parts separately
- Always evaluate two samples
- Weight sum always positive
- 2X the cost
- Low variance

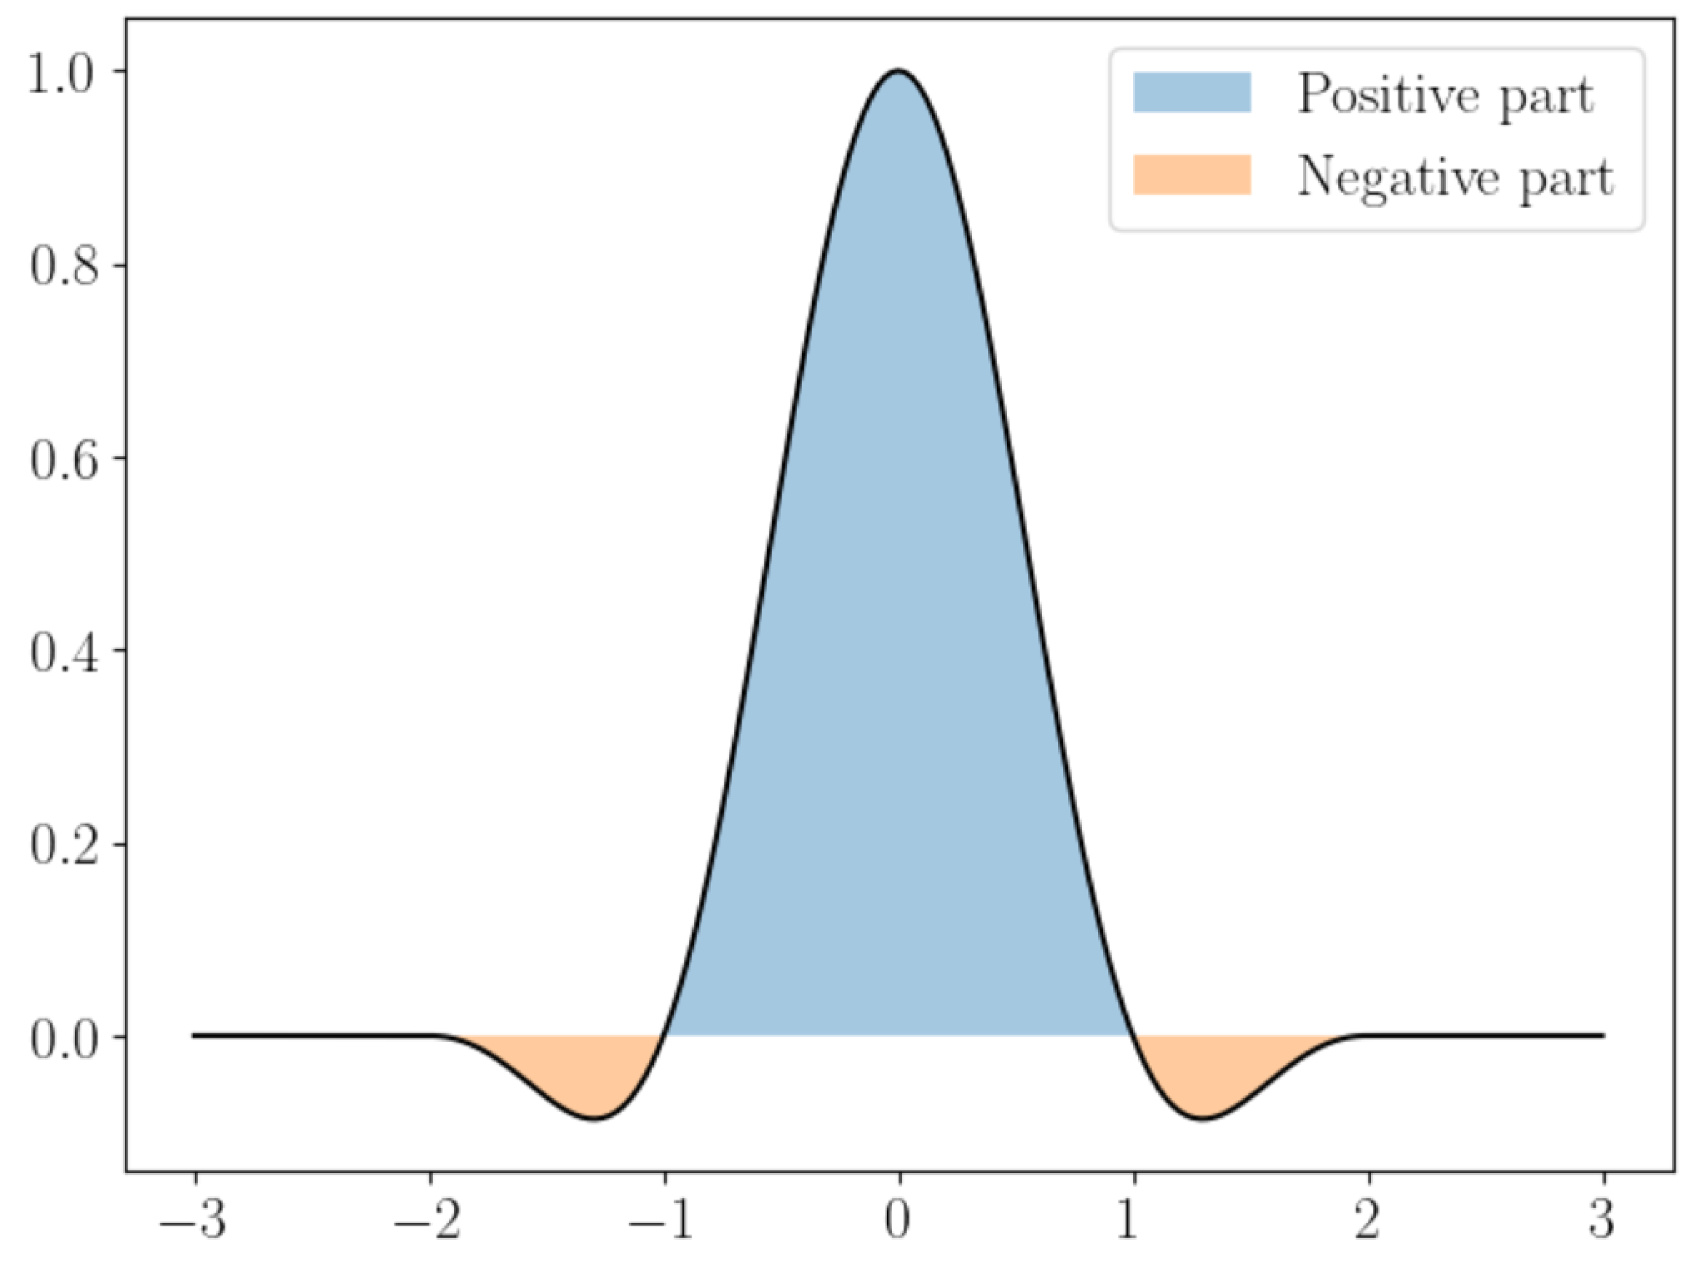

$$\langle F\rangle = \sum_i w_i^+ t_{j^+} - \sum_i w_i^- t_{j^-} .$$

Positivization – Results

Bilinear

Mitchell