Sampling and Nyquist's theorem

D.A. Forsyth

University of Illinois at Urbana Champaign

Sampling in 2D

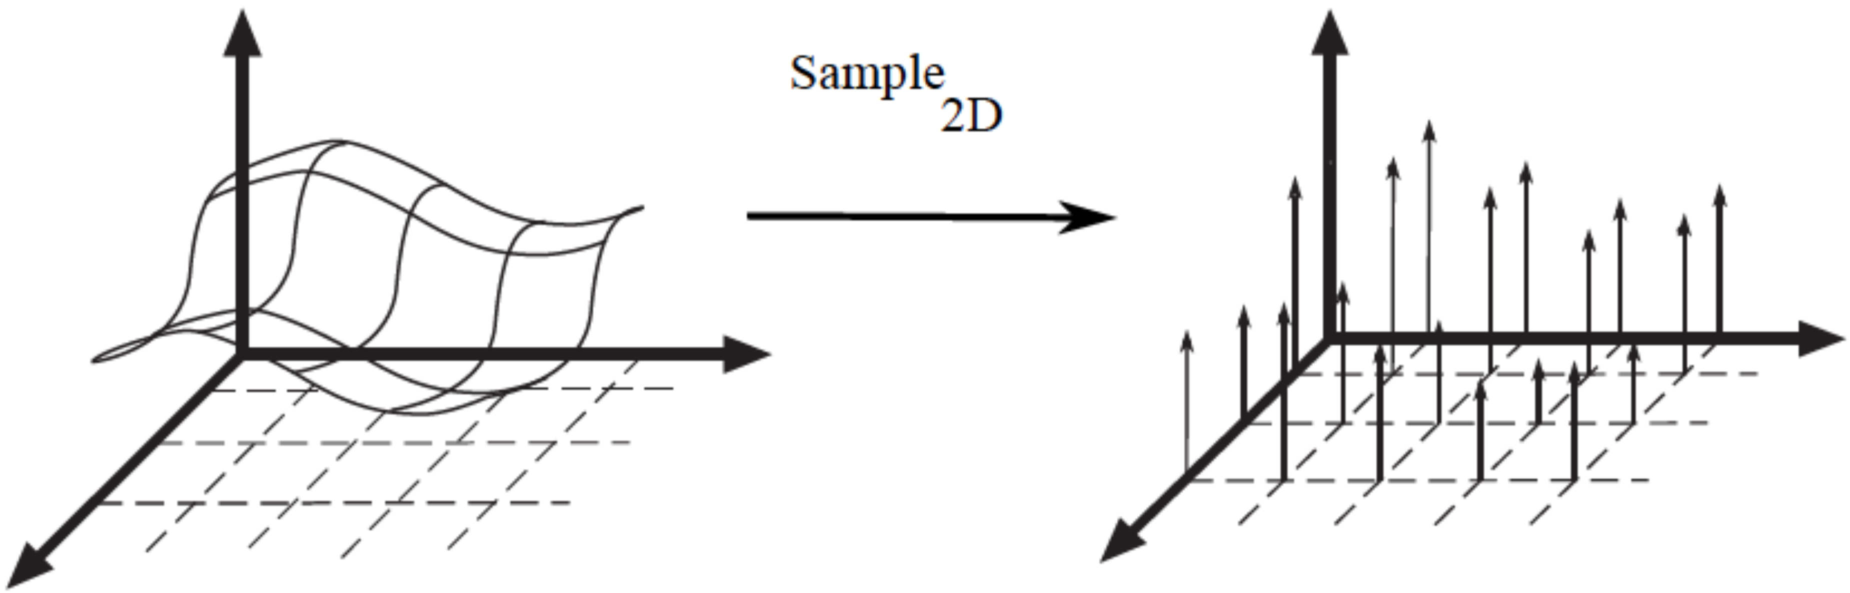

FIGURE 7.4: *Sampling in 2D takes a function and returns an array; again, we allow the array to be infinite dimensional and to have negative as well as positive indices.*

Modelling a sampled function

12.2.1 Modelling a Sampled Function

A crucial step is a reasonable model of a sampled function. Passing from a continuous function—like the irradiance at the back of a camera system—to a collection of values on a discrete grid —like the pixel values reported by a camera—is referred to as *sampling*. Sampling must lose information about the original function (for example, see Figure 3.4). Accounting for what is lost requires building a model of the sampling process quite carefully.

Write

$$\texttt{sample}_{2D}(f)$$

for an operation that takes a continuous function in 2D and returns a sampled version. The sampled version should represent the values of f at all integer points (you can get any other uniform grid with a scale). It is highly desirable that $\texttt{sample}_{2D}(f)$ produce a result that is compatible with integration. In particular, that

$$\int g(u,v)\texttt{sample}_{2D}(f)dudv \approx \int g(u,v)f(u,v)dudv$$

Modelling a sampled function

to the extent possible for any $g(u)$. Recall the definition of the δ function in 2D (Definition 11.3). It turns out that the right choice for $\texttt{sample}_{2D}(f)$ is

$$\texttt{sample}_{2D}(f) = \sum_{ij} f(i,j)\delta(x-i, y-j)$$

The grid is infinite in each dimension to avoid having to write ranges, etc. (Figure 12.2). The δ function is a conceptual device to make the mathematical plumbing work properly. There is no need to place one at each sample function in an array inside your programs (and you can't – you'd have to have an opinion about the value of $\delta(0)$, which isn't going to work out). This definition yields a model which behaves well for integrals. In particular,

$$\int g(u)\texttt{sample}_{2D}(f)du = \sum_{ij} f(i,j)g(i,j)$$

Interpolation is by convolution

Recall the interpolate of Section 3.1 had the form

$$\mathcal{I}(x,y) = \sum_{i,j} \mathcal{I}_{ij} b(x-i, y-j).$$

Here b is some function with the properties $b(0,0) = 1$ and $b(u,v) = 0$ for u and v any other grid point. This is linear and shift invariant (**exercises**) so it must be a convolution. The way to see the convolution is to use the model of sampling, above. This exposes the convolution in interpolation. Check that

$$\begin{aligned}
\texttt{sample}_{2D}(\mathcal{I}) * b &= \int\int \sum_{ij} \mathcal{I}_{ij} \delta(x-u-i, y-v-j) b(u,v) du dv \\
&= \sum_{ij} \mathcal{I}_{ij} \int\int \delta(x-u-i, y-v-j) b(u,v) du dv \\
&= \sum_{i,j} \mathcal{I}_{ij} b(x-i, y-j) \text{ from the property of a } \delta \text{ function}
\end{aligned}$$

The FT of a sampled signal

$$
\begin{aligned}
\mathcal{F}(\texttt{sample}_{2D}(f(x,y))) &= \mathcal{F}\left(f(x,y)\left\{\sum_{i=-\infty}^{\infty}\sum_{j=-\infty}^{\infty}\delta(x-i,y-j)\right\}\right) \\
&= \mathcal{F}(f(x,y)) * \mathcal{F}\left(\left\{\sum_{i=-\infty}^{\infty}\sum_{j=-\infty}^{\infty}\delta(x-i,y-j)\right\}\right) \\
&= \sum_{i=-\infty}^{\infty}\mathcal{F}(f)(u-i,v-j),
\end{aligned}
$$

The FT of a sampled signal

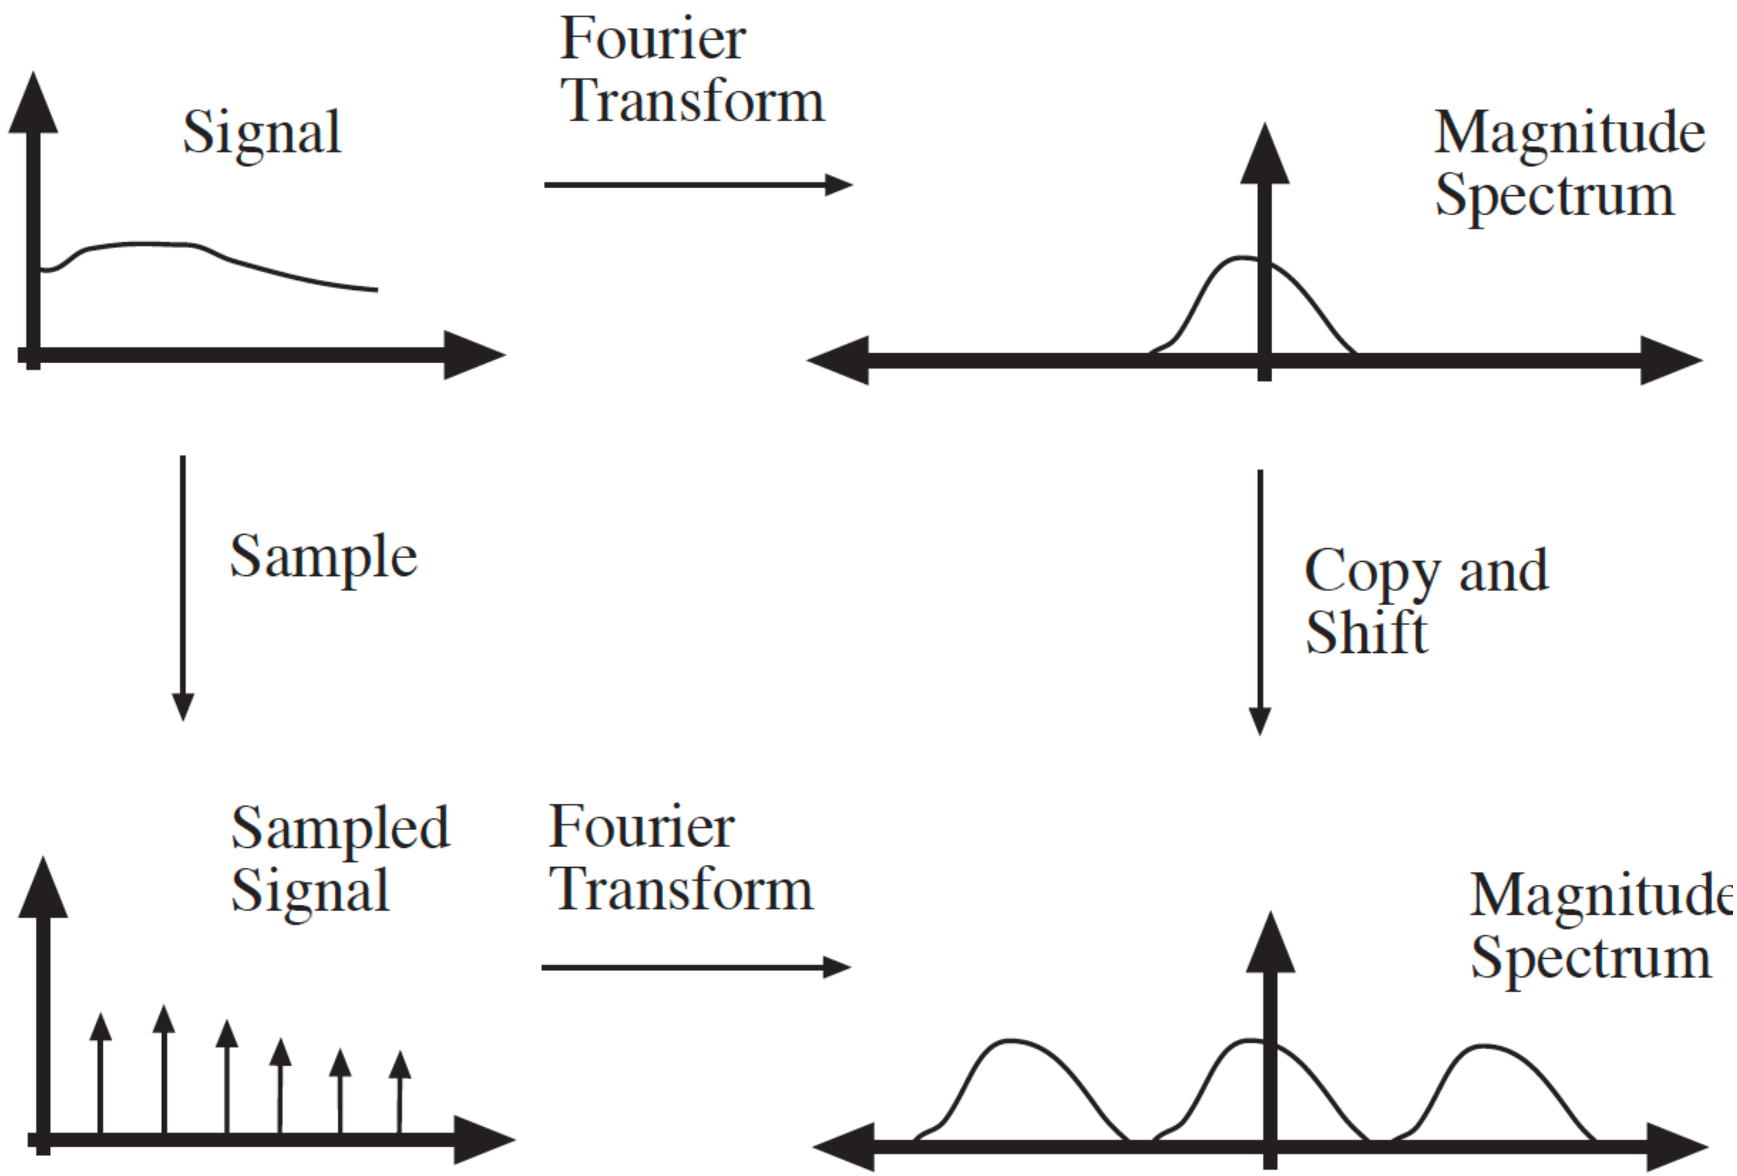

Sampling

- If the magnitude blobs don't overlap, you can reconstruct from samples
- Just cut out the blob in FT space with a box filter, inverse FT

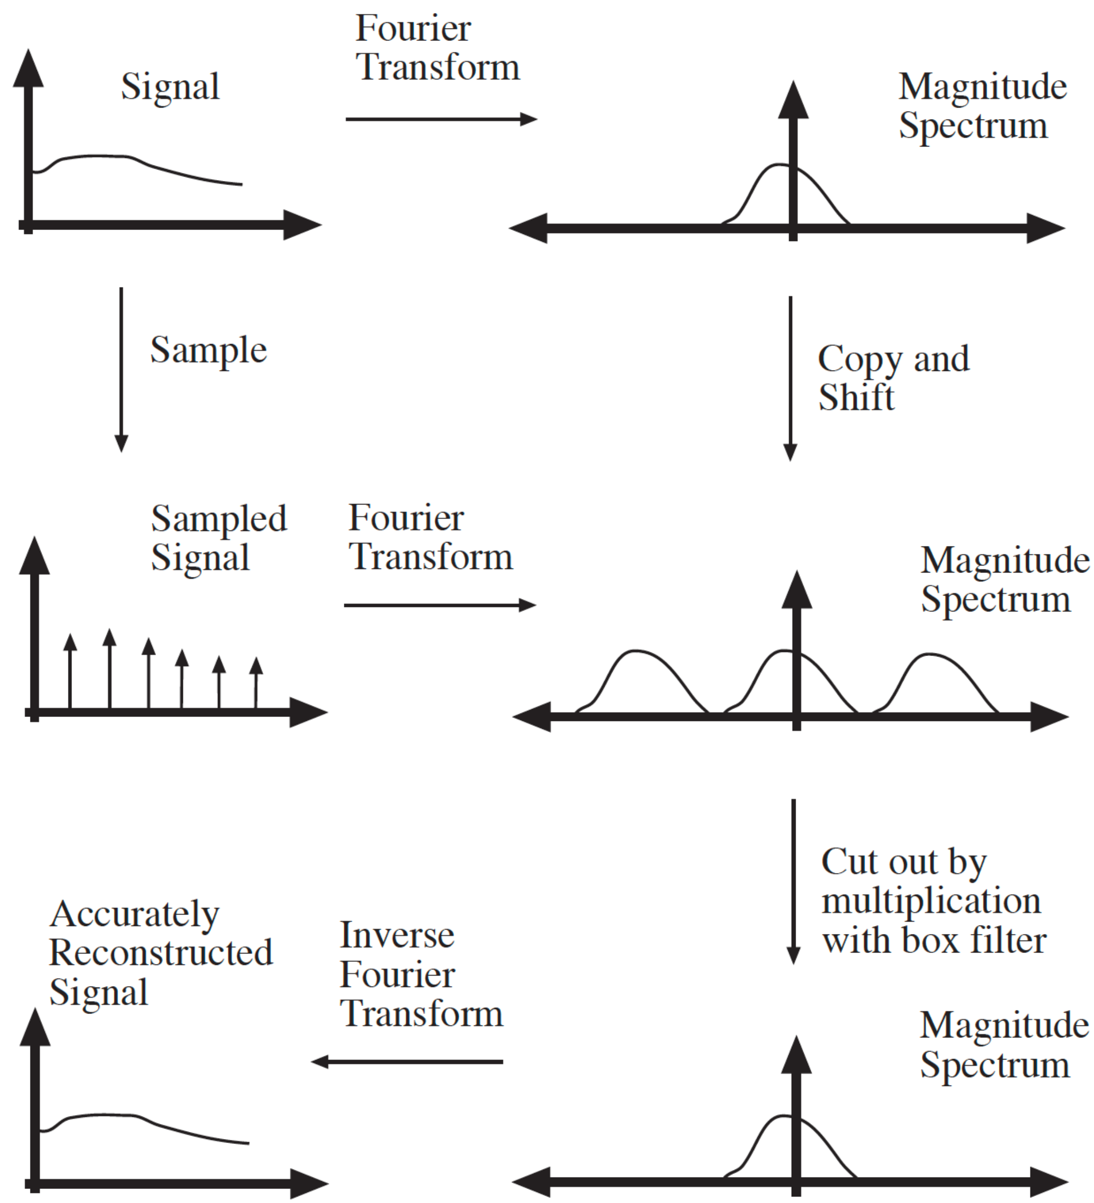

Sampling

- If the magnitude blobs overlap, you can't reconstruct from samples
- Cutting out won't work – you'll get some information from a different frequency that aliases
- Nyquist's theorem

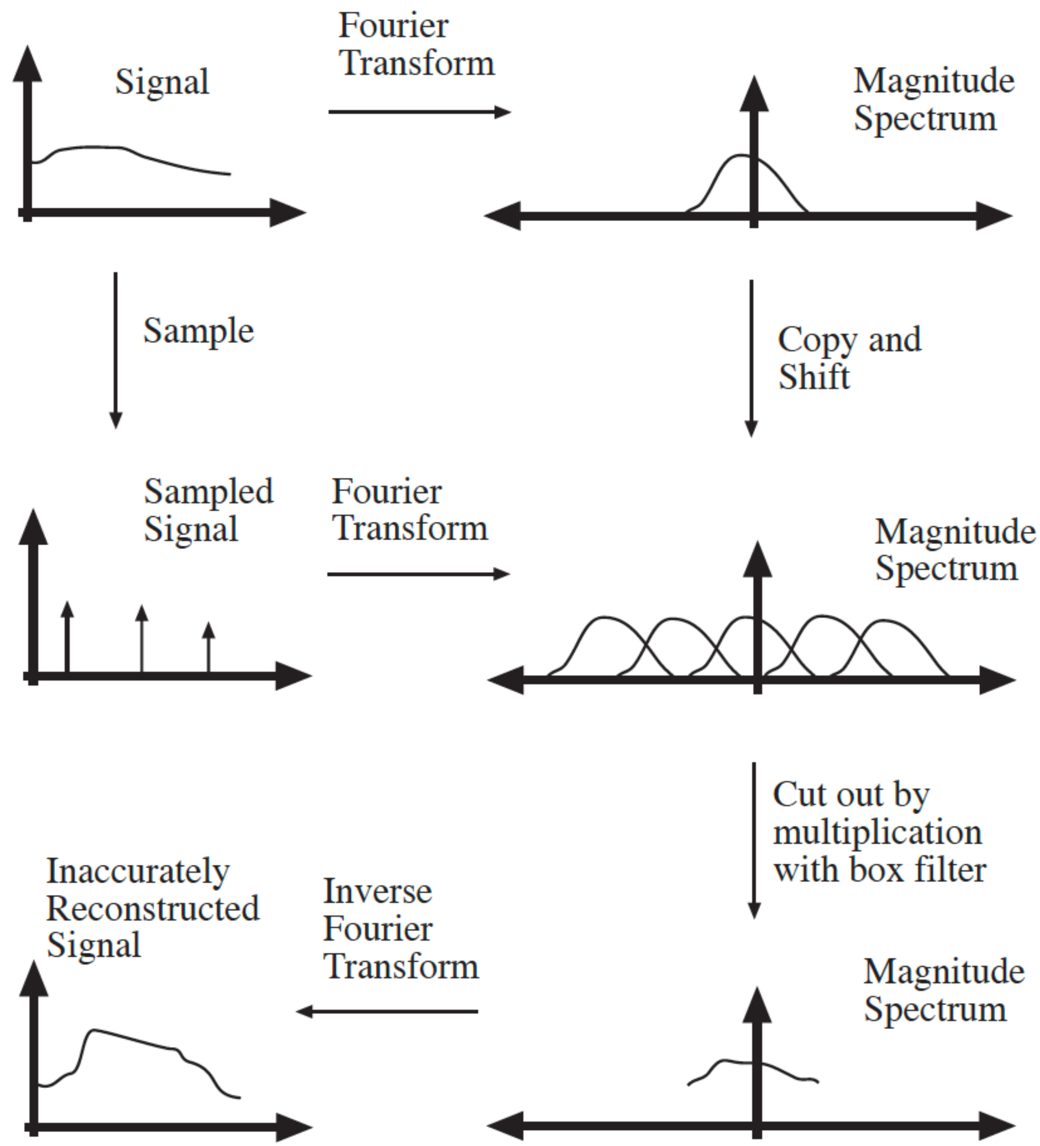

Consequences

- Nyquist limits aren't really viable
  - Apply the convolution theorem
  - A box in FT magnitude space is a filter with infinite support (and you can't make one of those)

- You're forced to choose a filter that is low pass, but isn't perfect
  - the choice has consequences
  - Gaussian is such a filter

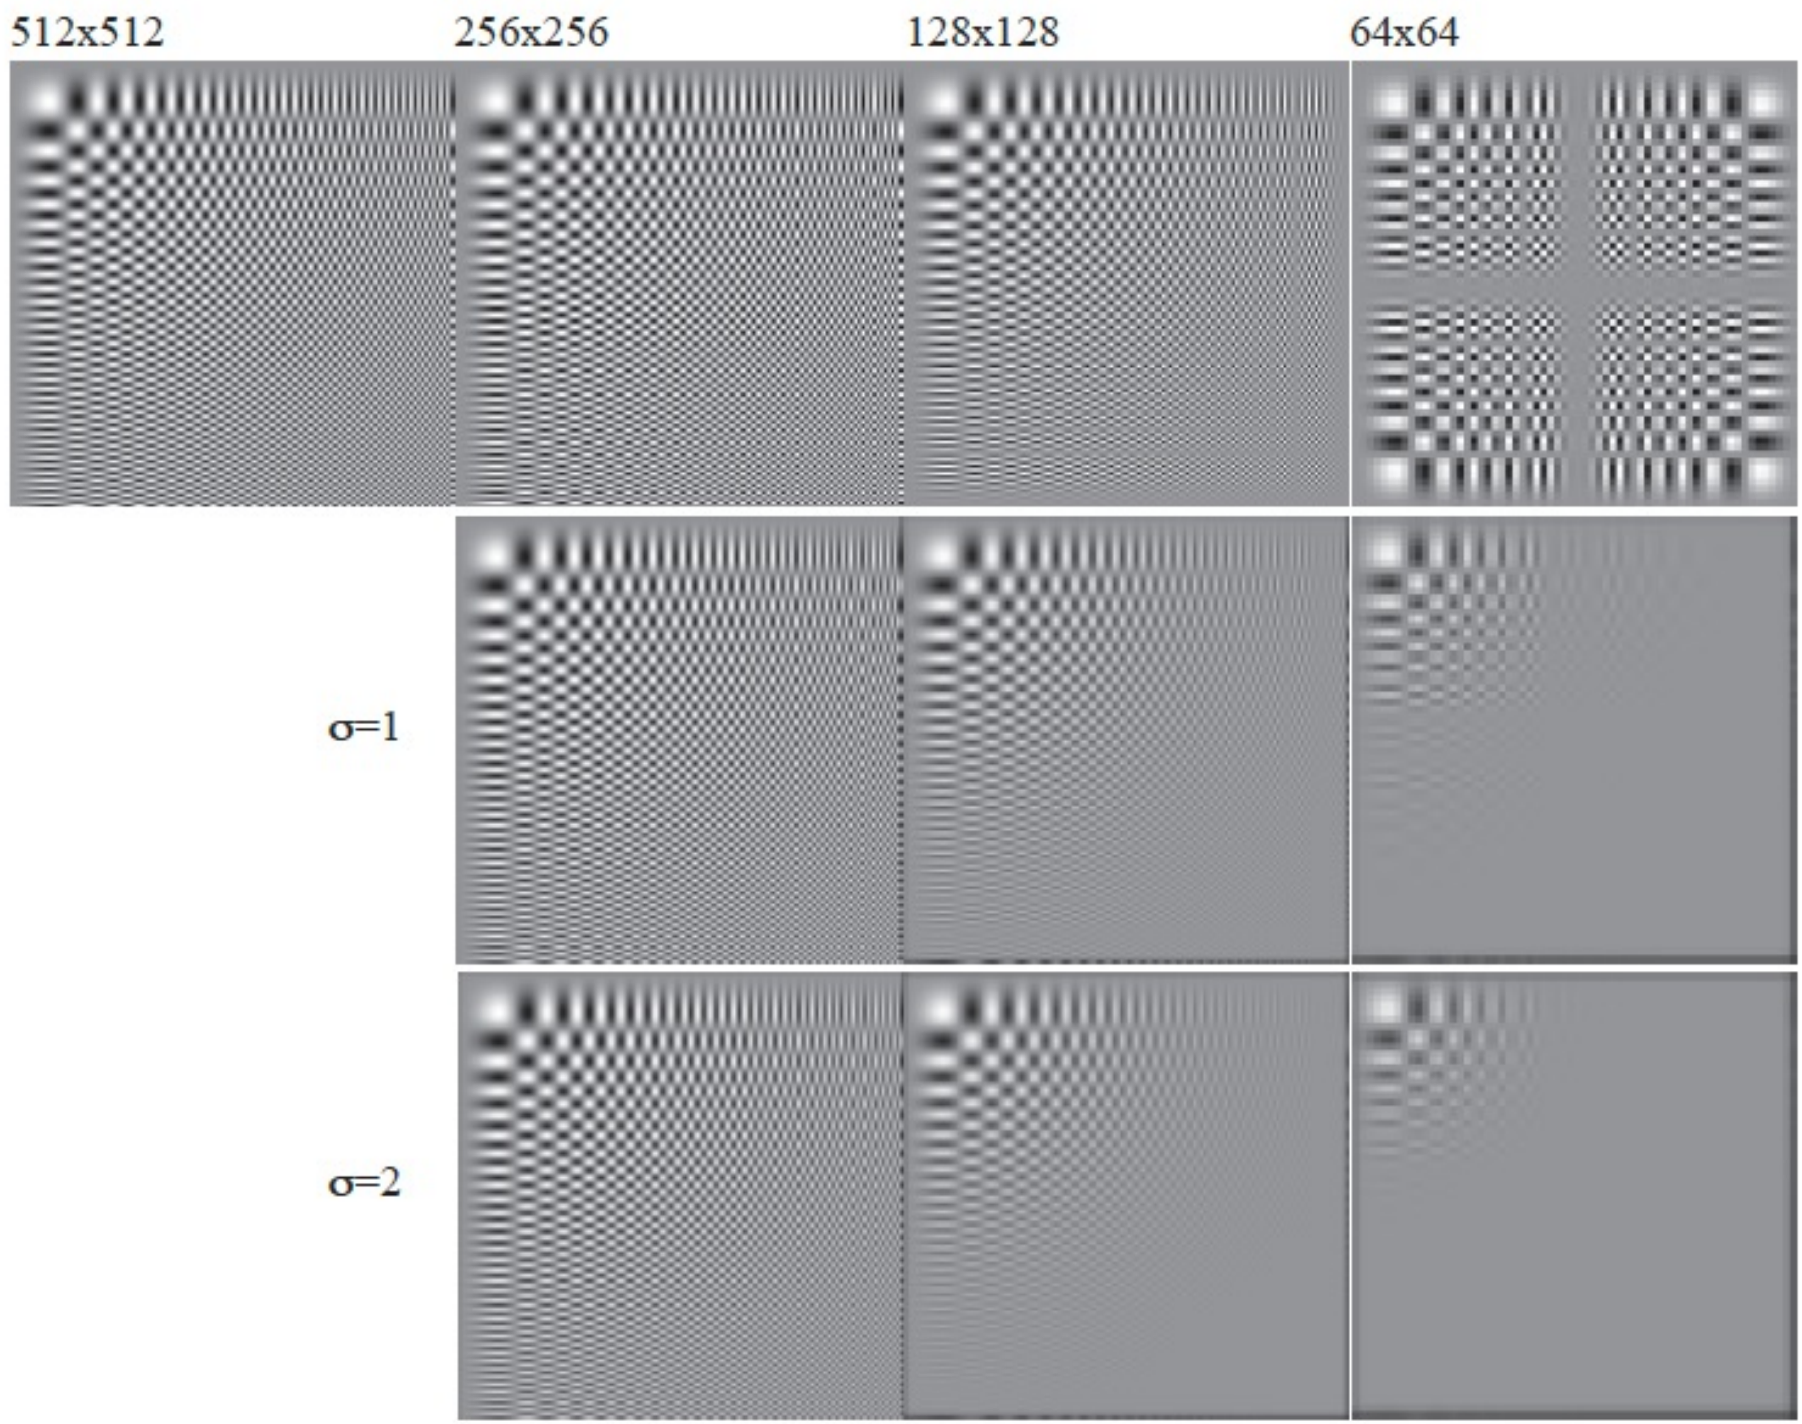
512x512
256x256
128x128
64x64
σ=1
σ=2

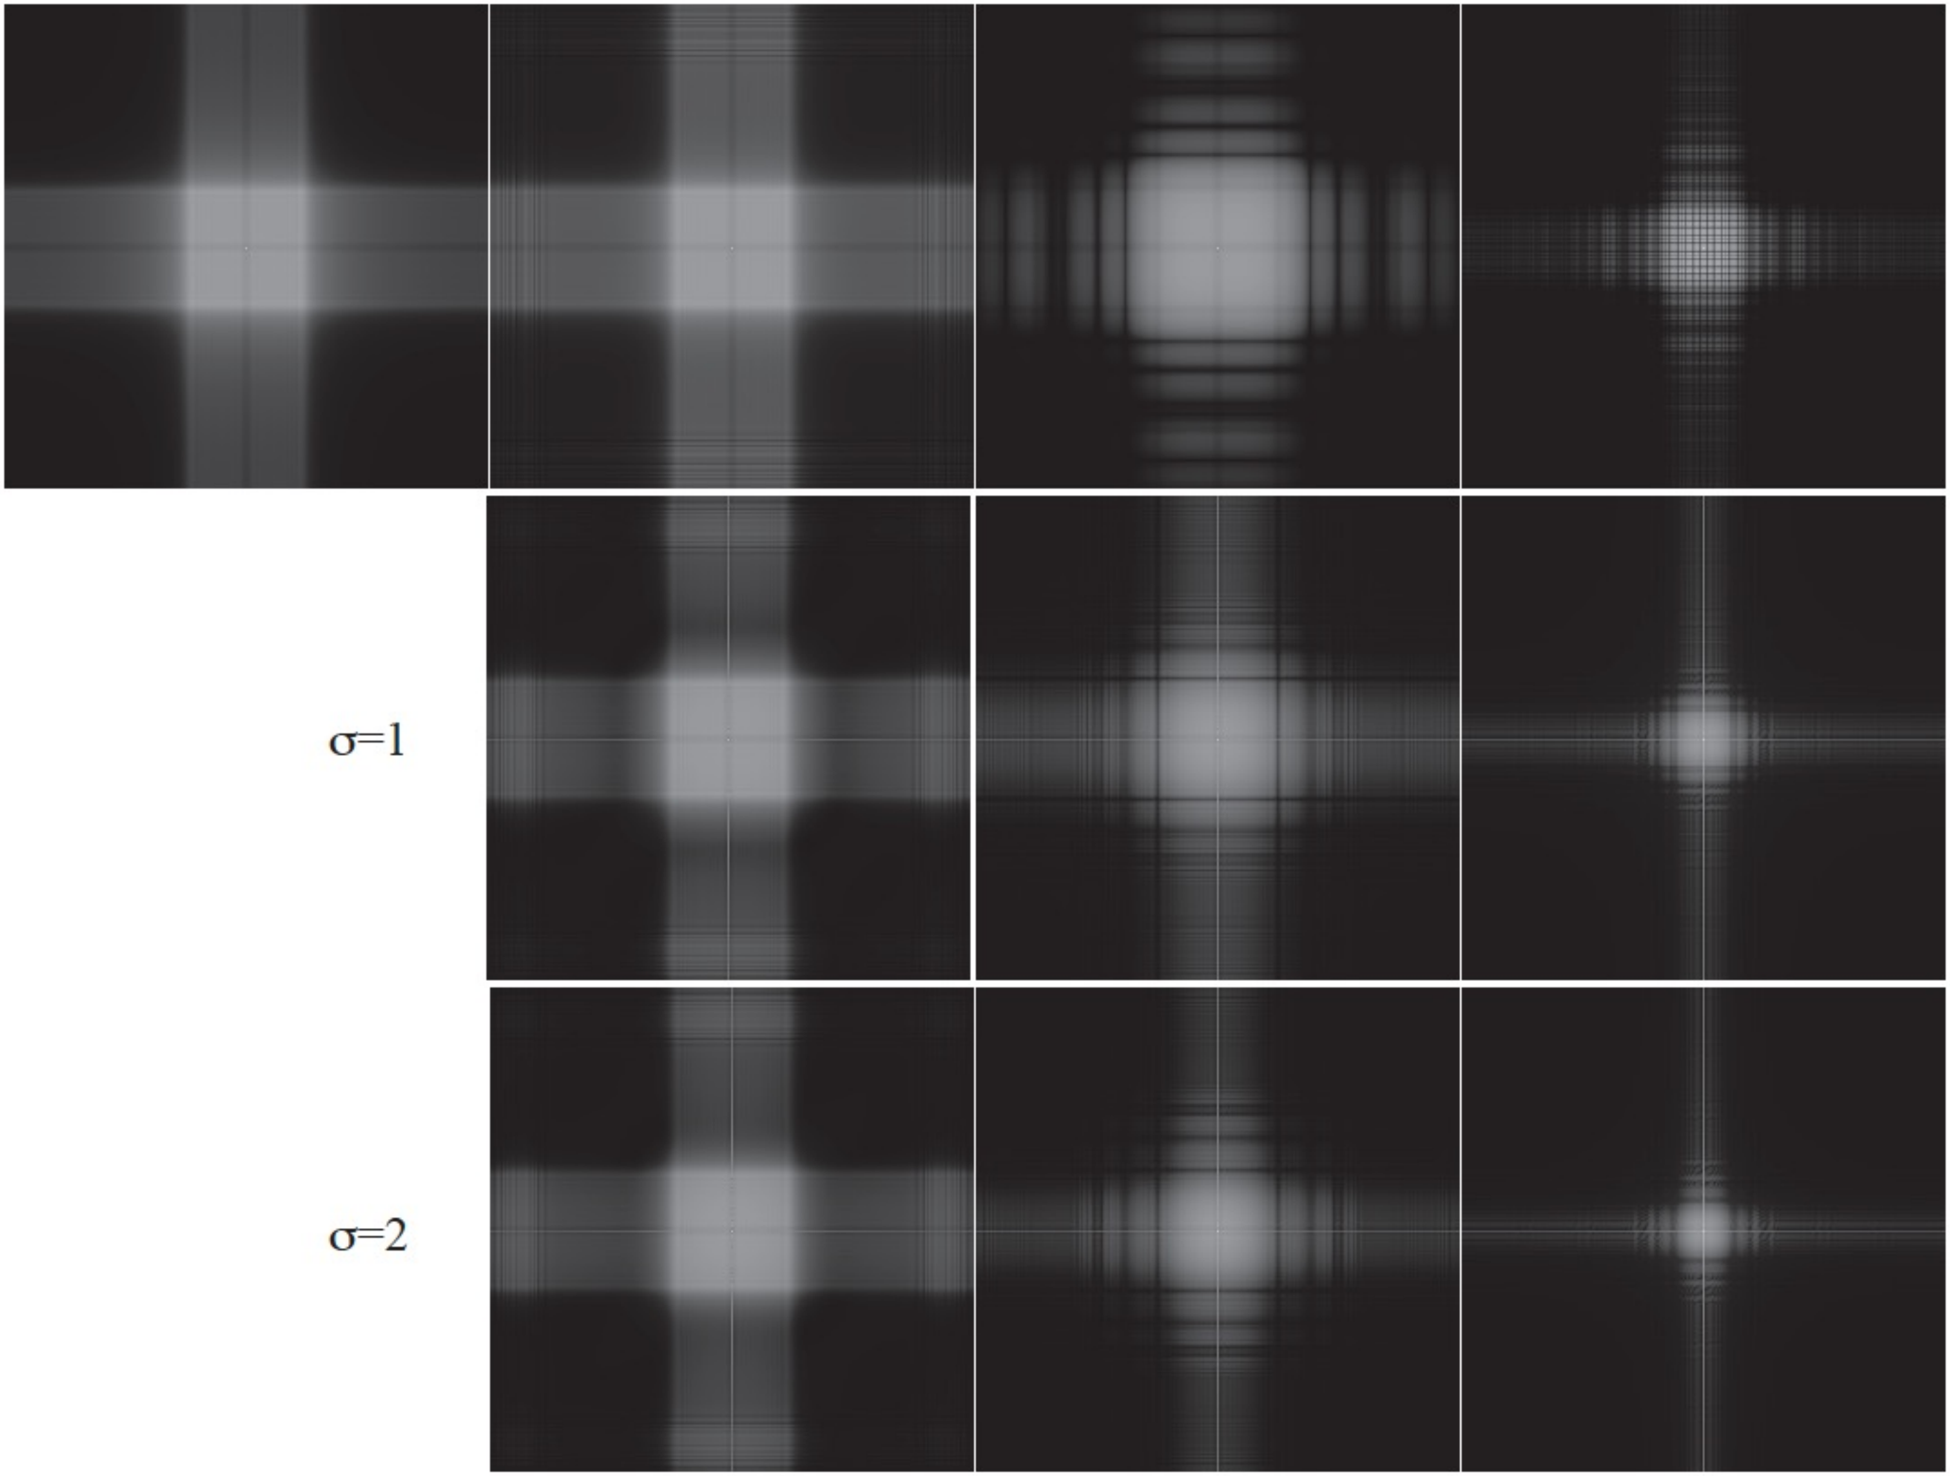
σ=1
σ=2

Mystery 2

- Why can downsampling sometimes lead to aliasing?

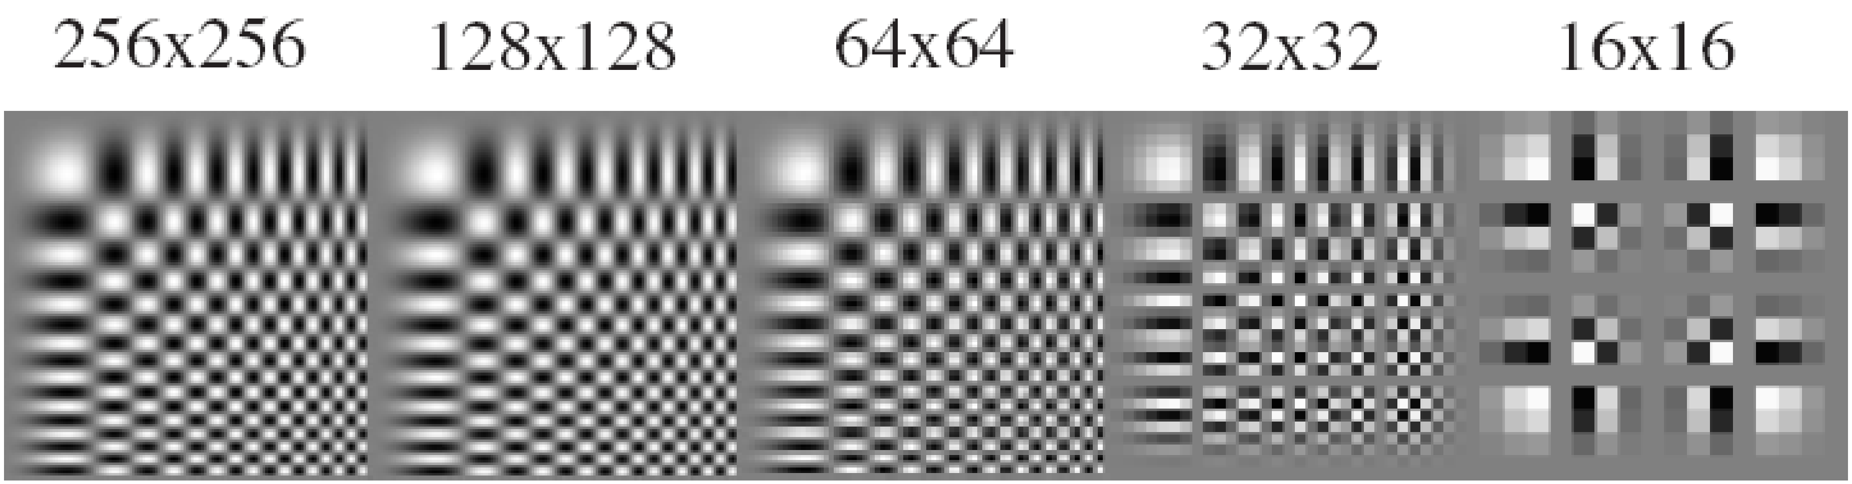

The downsampling mangles the Fourier Transform magnitude spectrum UNLESS you smooth by enough; and if you do, you lose some information

Mystery 2 SOLVED

- Why can downsampling sometimes lead to aliasing?

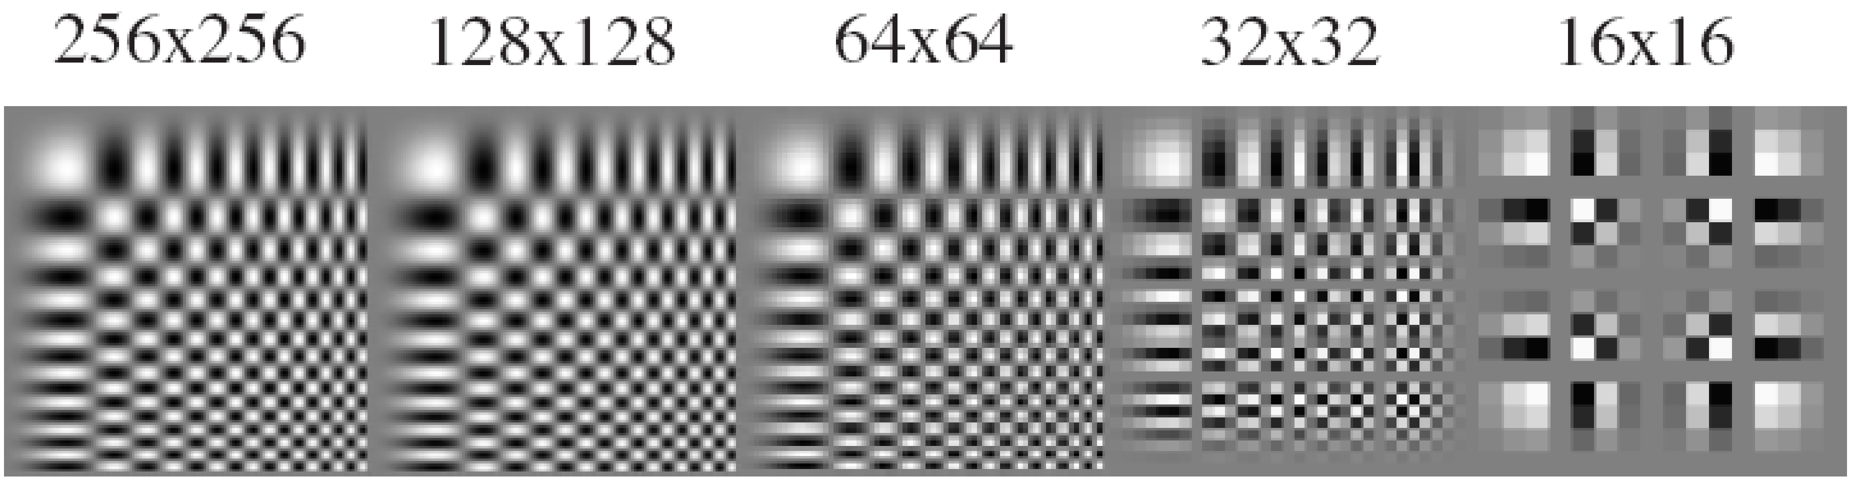

- The downsampling mangles the Fourier Transform magnitude spectrum UNLESS

Analyzing interpolation methods

- Perfect reconstruction requires convolution with a sinc filter in the spatial domain,
  - which is bad because sinc has infinite support
- Instead, simpler reconstruction (interpolation) methods are typically used

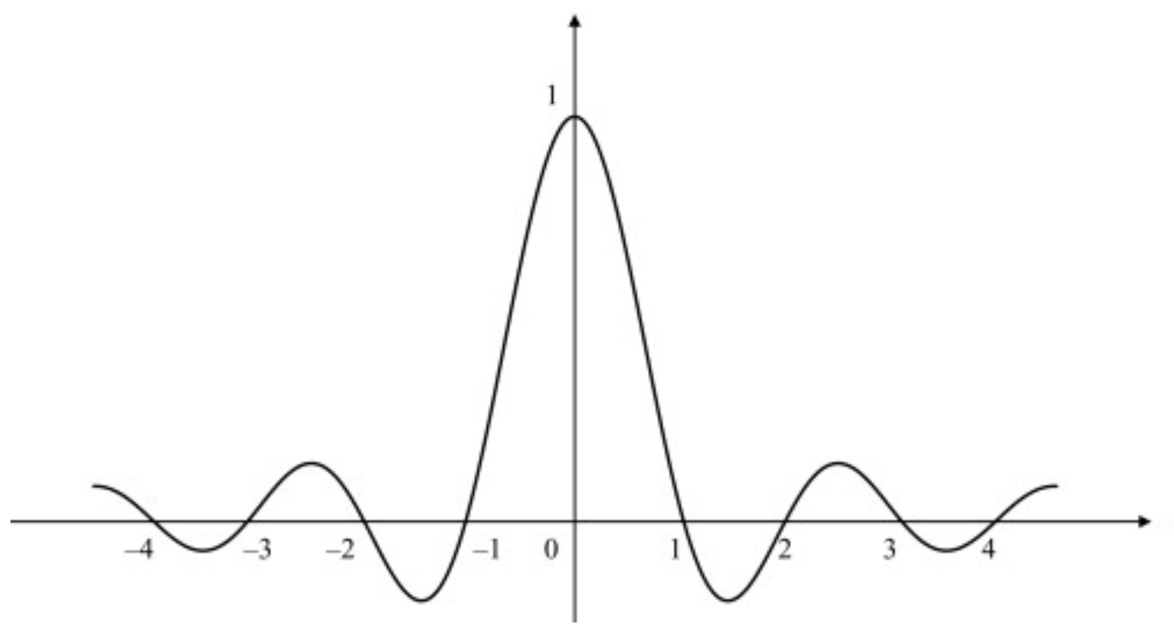

Analyzing interpolation methods

- Linear reconstruction: convolve with *triangle filter*:

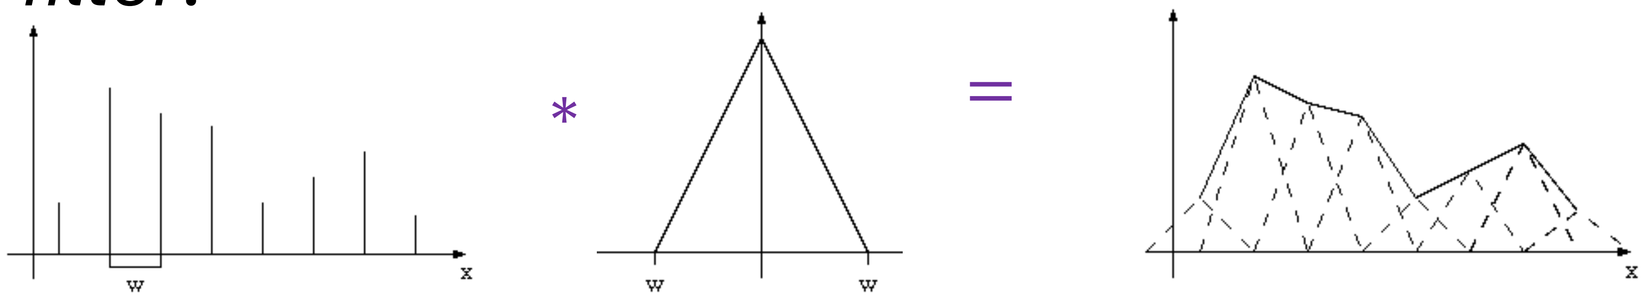

- Fourier transform of triangle filter is the sinc^2 function,
  - so multiplying the signal's spectrum by it introduces high-frequency artifacts

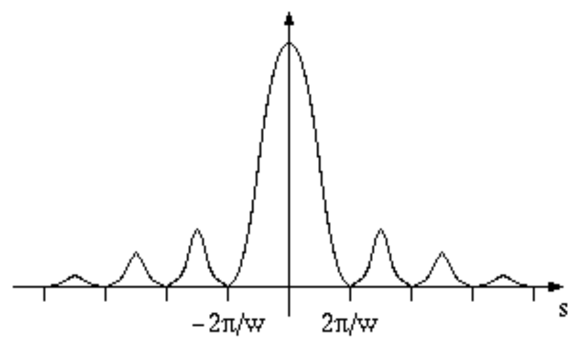

Image source

Bilinear interpolation closeup

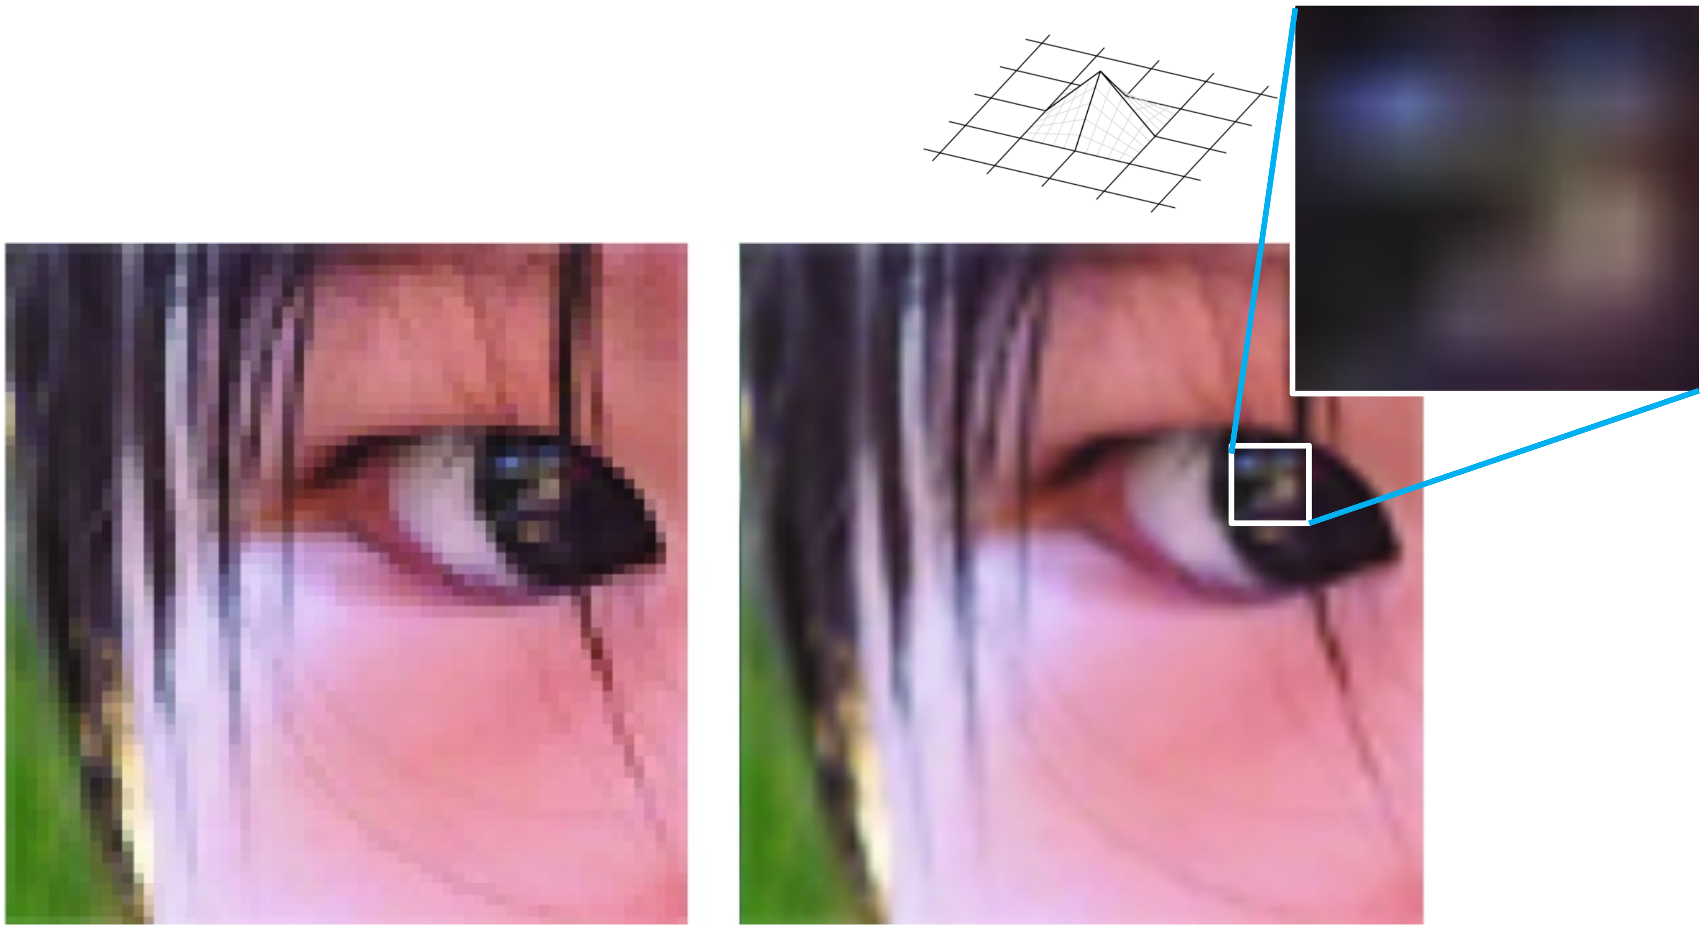

Image source

Why else should you care?

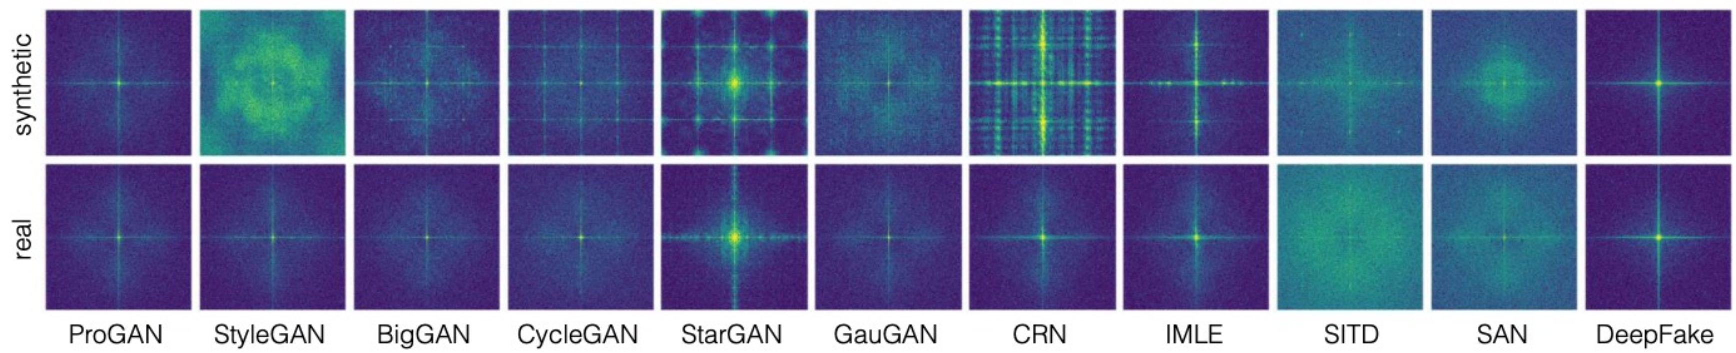

Figure 7: **Frequency analysis on each dataset.** We show the average spectra of each high-pass filtered image, for both the real and fake images, similar to Zhang *et al*. [50]. We observe periodic patterns (dots or lines) in most of the synthetic images, while BigGAN and ProGAN contains relatively few such artifacts.

S.-Y. Wang et al. CNN-generated images are surprisingly easy to spot… for now. CVPR 2020

Why else should you care?

Checkerboard and repetition artifacts in GAN-generated images

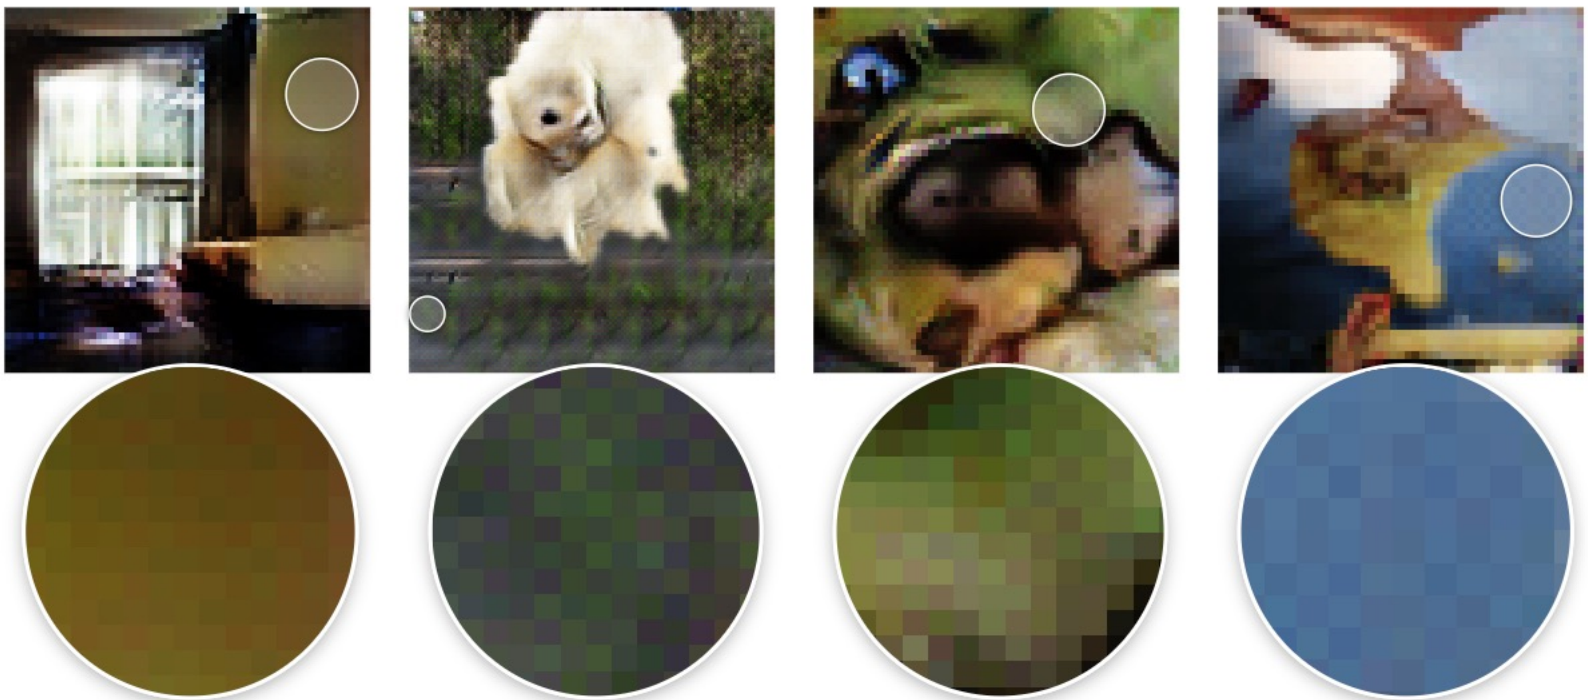

Radford, et al., 2015 [1] Salimans et al., 2016 [2] Donahue, et al., 2016 [3] Dumoulin, et al., 2016 [4]

https://distill.pub/2016/deconv-checkerboard/

Things to think about...

12.7. Check that $\int g(u)\texttt{sample}_{2D}(f)du = \sum_{ij} f(i,j)g(i,j)$.

12.8. Section 22.6 has:"This definition yields a model which behaves well for integrals. In particular,

$$\int g(u)\texttt{sample}_{2D}(f)du = \sum_{ij} f(i,j)g(i,j)$$

which is the best approximation to the integral that you will get if you know $f(u,v)$ only at integer points." Explain.

12.9. Write an expression for what you would get if you convolve $\texttt{sample}_{2D}(\mathcal{I})$ with $g(x,y)$, then sample the result.

12.10. Write an expression for what you would get if you convolve $\mathcal{I}$ with $\texttt{sample}_{2D}(g)$, then sample the result.

12.11. Section 12.2.3 has:"convolving a function with a shifted δ-function merely shifts the function". Show this is true.

12.12. Section 12.2.4 has: "If the sampled image is downsampled by two, for example, the copies now have centers on the half-integer points in u, v space." Explain.